MIT Technology Review

Tech's 50 Smartest Companies

VOL. 120 NO. 4 JULY/AUGUST 2017 US $6.99/CAN $7.99



The Business Issue


Build more than a network.
Build an earth-shattering,
gravity-defying path to the future.
We believe in more than routers and switches. We believe in igniting possibility and challenging convention.
See how we can help tell your story at Juniper.net/BuildMore

JUNIPER®
NETWORKS

The Business Issue



Nanette Byrnes is* MIT Technology Review*'s senior editor for business.

What does it mean to be on our list of the 50 Smartest Companies?

It does not necessarily mean that you are the biggest in your industry—though we have some of those—or even the most profitable. We don't depend on a tally of patents granted. Nor do we care where a company operates or how large its market capitalization is.

Instead, the list that you will find on pages 56 and 57 is our take on two factors: technological leadership and business acumen. These 50 companies combine a high level of technology innovation with a business model that will help them make the most of it. Public and private, large and small, based in countries around the globe, this group of companies is creating new opportunities and pouncing on them. These are the ones that competitors must follow.

Artificial intelligence, a technology that we believe will drive much of the economic growth over the next few years, is sprinkled throughout the list. The company at the top, Nvidia, has gained expertise in AI and used it to transform itself. Once known as a maker of chips for gaming, it is now a leading player in deep learning and autonomous vehicles. Amazon, No. 3, is on the list again for its ambitions to build an AI-powered store and place the technology at the heart of the home of the future.

Companies focused on DNA analysis and developing gene therapies are a dominant group on the list—23andMe (No. 4), Spark Therapeutics (No. 10), Illumina (No. 22), and Oxford Nanopore (No. 32) among them.

Since 2010, SpaceX has been highlighted six times. It's on again in 2017 (at No. 2) on the strength of its ability to stick rocket landings and then recycle its crafts for another go-round, potentially changing the economics of space travel. IBM has made the list every one of the past seven years; this year it's at No. 39 for its work with blockchain, cloud AI, and quantum computing.

Many companies are here for the first time. Among them are companies like China's iFlytek, debuting at No. 6, which has created voice-assistant software that can accurately translate between Chinese and a number of other languages (including English, German, and Uighur), overcoming challenges such as dialect, slang, and background noise. Another new name is Face ++, No. 11, a Chinese startup. Face ++ is using its novel face-recognition technology to help detect fraud and make payments more convenient.

It's been a tough year even for some of our smartest companies. Tesla, No. 4 last year, fell to 31 this year, hurt by maintenance problems, issues with its Autopilot feature, and the challenges of absorbing the nation's largest installer of solar panels into its business following its $2.6 billion purchase of SolarCity.

We also decided that some well-known firms just aren't smart enough these days. Neither Uber nor Airbnb made the list. In the case of Uber, a cascade of bad news about its business practices, management style, and claimed technology theft combined to keep it off. For Airbnb, it wasn't a matter of bad news but, rather, a lack of big technology news that cost it a slot among our top 50.

Seen all together, as our editor David Rotman argues in his essay introducing the list ("It Pays to Be Smart," page 54), this is our best guess as to the firms that will be the dominant companies of the future. These are the companies with the inside track on the technologies, such as artificial intelligence, that will define the economy in the coming years—and have the best shot at building great businesses around them.

ROXANA PERDUE



Merck invents to
Keep Love Alive

Take Everett and Evelyn. Healthy. Moving forward together. And still as much in love as they were when they met in 1944—back when he was a paratrooper and she was a nurse.

For more than a century, Merck has been inventing medicines and vaccines for many of the world's most challenging diseases. Today, we're inventing new ways to treat cancer, with more than 450 clinical studies underway. All in an effort to help keep people healthy, hopeful, and allow their relationships to endure for years to come.

Learn more at Merck.com/InventingForLife

MERCK
INVENTING FOR LIFE

Copyright ©2017 Merck Sharp & Dohme Corp., a subsidiary of **Merck & Co., Inc.,** Kenilworth, NJ USA. All Rights Reserved. CORP-1210605-0007 06/17

The Business Issue 2017

Table of Contents

July/August 2017

2 **From the Editor**

8 **Feedback**

VIEWS

10 **Virtual Trolls**
Anonymity is stifling social VR's tremendous potential.

10 **The Innovator Gap**
Our innovation ecosystem is failing our best technologists.

12 **Drones to the Rescue**
Using UAVs to deliver drugs is worth every penny.

UPFRONT

15 **Reinventing Rice**
Climate change is coming. We'll need hardier crops.

20 **Who's Winning at AI?**
The story is told in the number of research papers published.

22 **Google's Quantum Leap**
With a new chip, it's aiming for "quantum supremacy."

23 **CRISPR That You Eat**
The gene-editing technology could replace antibiotics.

24 **RNA Miracle Drug**
One boy's revival may point the way to the future of medicine.

ON THE COVER



Illustration by David Biskup

26 **Is It Real? Only AI Knows**
It's getting easier to generate realistic fake audio and video.

27 **The Laser Problem**
Lidar sensors are holding up progress with self-driving cars.

Q+A

34 **A Job Plan for Humans and Robots Alike**
Melonee Wise talks about her plans for Fetch Robotics.

52 **Search Giant Baidu's Growing Ambitions**
Ya-Qin Zhang thinks AI will turn his company into a global powerhouse.

REVIEWS

84 **Virtual Reality's Missing Piece: People**
The real promise of VR may be as a new way to communicate.
By Rachel Metz

88 **The Unaffordable Paradise**
Tech startups saved the cities. Then wrecked them.
By Richard Florida

92 **Acquion Energy**
Why is it so hard for energy startups to thrive?
By James Temple

96 **Our Biased Algorithms**
Machines now make big decisions about our lives. Can they do it fairly?
By Matthias Spielkamp

40 YEARS AGO

100 **The Inevitable EV**
Four decades ago, electric cars looked ready to take over the daily commute.

FEATURES

Why Tesla Is Worth More Than GM p28



A few companies mastering digital technology now own huge chunks of the economy. Does that explain slow growth? *By James Surowiecki*

Help from Above p36



In Rwanda, an early commercial test of unmanned aerial vehicles is delivering urgently needed medical supplies. *By Jonathan W. Rosen*

50 Smartest Companies p54-83



It Pays to Be Smart *by David Rotman* p54

The List *by the Editors* p56

Alphabet *by Tom Simonite* p58

Carbon *by Katherine Bourzac* p62

Oxford Nanopore *by Antonio Regalado* p66

IBM *by David Freedman* p72

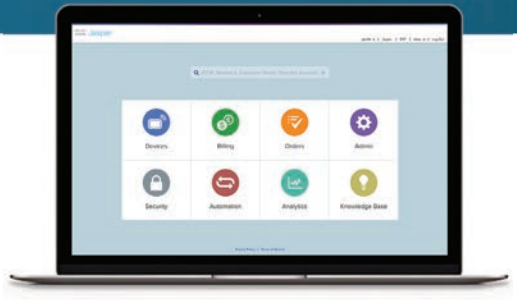
General Electric *by Elizabeth Woyke* p78

EASILY CRUNCH DATA FROM MILLIONS OF CONNECTED DEVICES AND
MAKE THAT DATA WORK FOR YOU

AUTOMATE AND SCALE
YOUR IoT SERVICES ACROSS THE GLOBE

PROTECT YOUR IoT DEVICES
AND YOUR BUSINESS FROM SECURITY THREATS

YOU'VE GOT IoT *UNDER CONTROL.*



cisco Jasper | Control Center 7.0

Announcing the New

ControlCenter7.com

EDITORIAL

Editor David Rotman, **Editor at Large** Brian Bergstein, **Deputy Editor** Megan Barnett, **Senior Editor, Business** Nanette Byrnes, **Senior Editor, MIT News** Alice Dragoon, **Senior Editor, AI and Robotics** Will Knight, **Senior Editor, Mobile** Rachel Metz, **Senior Editor, Biomedicine** Antonio Regalado, **Senior Editor, News and Commentary** Michael Reilly, **Senior Editor, Energy** James Temple, **San Francisco Bureau Chief** Tom Simonite, **Executive Producer** Kyanna Sutton, **Business Editor** Elizabeth Woyke, **Managing Editor** Timothy Maher, **Copy Chief** Linda Lowenthal, **News and Commentary Editor** Jamie Condliffe, **Associate Editors** Emily Mullin, Mike Orcutt, **Social Media Editor** Julia Sklar, **Senior Production Director** James LaBelle, **Contributing Editors** George Anders, Katherine Bourzac, Peter Burrows, Jon Cohen, Peter Fairley, Simson L. Garfinkel, Courtney Humphries, Amanda Schaffer, David Talbot

DESIGN

Chief Creative Officer Jordan Awan, **Art Director** Sam Jayne, **Designer** Lynne Carty, **Art Assistant** Bryan Fountain

PRODUCT DEVELOPMENT

Chief Digital Officer and VP, Product Development Erik Pelletier, **Product Manager** Vanessa DeCollibus, **User Interface/Digital Designers** Emily Waggoner, Jon Akland, **Engineers** Shaun Calhoun, Molly Frey, Jason Lewicki, DJ Hartman, Zach Green

EVENTS

SVP, Events and Strategic Partnerships Amy Lammers, **Director of Events Programming** Laura Janes Wilson, **Senior Events Coördinator** Nicole Silva

SOLVE

Executive Director Alex Amouyel, **Managing Director, Community** Dawda Jobarteh, **Community Directors** Hala Hanna, Pooja Wagh, **Events Program Director** Matthew Minor, **Partnerships Manager** Bruna Braga, **Online Engagement Manager** Ken Meyer, **Community Relations Senior Officer** Alexander Dale, **Financial and Reporting Officer** J. Richard Crowley III, **Community Relations Officers** Patrick Diamond, Caroline Federal, Sara Monteabaro, **Program Operations Coördinator** Rachael Drew, **Speaker Relations Coördinator** Noya Leve

CORPORATE

Chief Executive Officer and Publisher Elizabeth Bramson-Boudreau

VP, Licensing and Communities Antoinette Matthews, **VP, International Business** Nicola Crepaldi, **Human Resources Manager** Hilary Siegel, **Assistant to the Chief Executive Officer** Katie McLean, **Manager of Information Technology** Colby Wheeler, **Office Manager** Linda Cardinal

FINANCE

Manager of Accounting and Finance Enejda Xheblati, **General Ledger Manager** Olivia Male, **Accountant** Letitia Trecartin

CONSUMER MARKETING

SVP, Consumer Revenues and Marketing Doreen Adger, **Director of Marketing and Communications** David W.M. Sweeney, **Senior Marketing Associate** Julie Swanson, **Assistant Consumer Marketing Manager** Katya Hill, **Director of Analytics** Tom Russell

ADVERTISING SALES

VP, Sales Laurie Hironaka / laurie@technologyreview.com / 415-640-5141, **Northwest Advertising Director** Marii Sebahar / marii@technologyreview.com / 415-416-9140, **Midwest Advertising Director** Paul Kissane / paul.kissane@technologyreview.com / 312-485-3840, **Business Development Manager** Debbie Hanley / debbie.hanley@technologyreview.com / 214-282-2727, **New York and Southeast Advertising Director** Ian Keller / ian.keller@technologyreview.com / 203-858-3396, **Northeast Advertising Director** Mason Wells / mason.wells@technologyreview.com / 917-656-2899, **Digital Sales Strategy Manager** Ken Collina / ken.collina@technologyreview.com / 617-475-8004, **Custom Editor** Mindy Blodgett, **Sales and Marketing Coördinator** Anna Raborn, **Advertising Services** webcreative@technologyreview.com / 617-475-8004, **Media Kit** www.technologyreview.com/media, **Germany** Michael Hanke / michael.hanke@heise.de / +49-511-5352-167, **China** Vincent Chen / +86-185-1033-0513, **Japan** Yoshimi Suezawa / adsales@technologyreview.jp / +81-3-3583-5364, **Spain and South America** Cecilia Nicolini / cecilia.nicolini@opinno.com / +34607720179

MIT ENTERPRISE FORUM, INC.

President Elizabeth Bramson-Boudreau, **Executive Director and Clerk** Antoinette Matthews, **Treasurer** Enejda Xheblati, **Director of Chapter Leadership and Process** Gaylee Duncan, **Director of Communications** Joyce Chen, **Marketing and Events Coördinator** Stephanie Eng

BOARD OF DIRECTORS

Martin A. Schmidt
Judith M. Cole
Jerome I. Friedman
Joichi Ito
Israel Ruiz
David Schmittlein
Alan Spoon

CUSTOMER SERVICE AND SUBSCRIPTION INQUIRIES

National
800-877-5230
International
903-636-1115
E-mail
customer_service@mittechnologyreview.info
Web
www.technologyreview.com/customerservice
MIT Records (alums only)
617-253-8270
Reprints
techreview@wrightsmedia.com
877-652-5295
Licensing and permissions
licensing@technologyreview.com

Technology Review
One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to equip its audiences with the intelligence to understand and contribute to a world shaped by technology.

Technology Review, Inc., is an independent nonprofit 501(c)(3) corporation wholly owned by MIT; the views expressed in our publications and at our events are not always shared by the Institute.

De Technologia non multum scimus. Scimus autem, quid nobis placeat.

Be a Part of the **Solution**



Solve is a community of leaders and change makers committed to solving the world's most pressing challenges through open innovation and partnership.



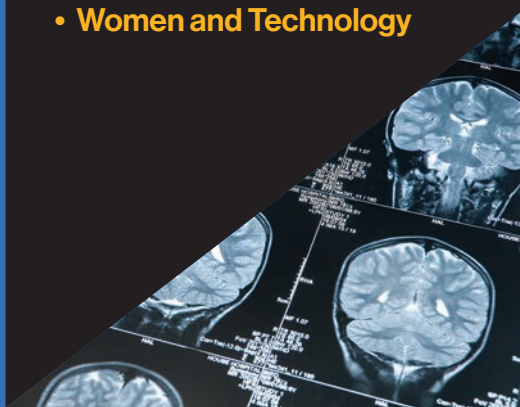
We're accepting solutions to four challenges until August 1, 2017:

- **Youth, Skills, and the Workforce of the Future**
- **Brain Health**
- **Sustainable Urban Communities**
- **Women and Technology**

Submit a solution:
solve.mit.edu/challenges

For membership inquiries:
membership@solve.mit.edu

SOLVE



Feedback

E-mail letters@technologyreview.com

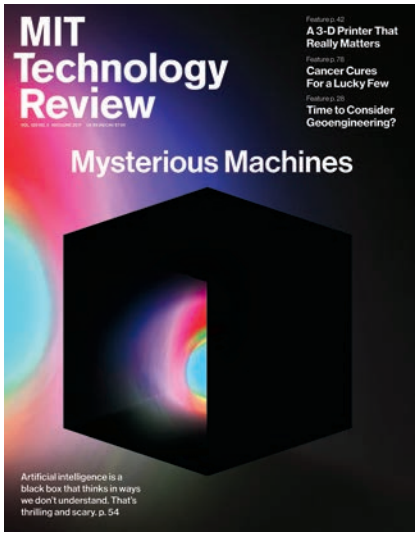
Write *MIT Technology Review*
One Main Street, 13th Floor
Cambridge, MA 02142

Please include your address, telephone number, and e-mail address.

Letters and comments may be edited for both clarity and length.

Letters and Comments

MIT Technology Review
Volume 120, Number 3



Of Course Complexity Is Baffling

In "The Dark Secret at the Heart of AI" (May/June 2017), Will Knight writes: "The workings of any machine-learning technology are inherently more opaque, even to computer scientists, than a hand-coded system." While it's partly true that hand-coded systems can be easier to understand than data-trained models, the issue isn't really that particular models are hard but rather that complex domains in general are hard to understand.

For example, do we understand why a certain stock goes up or down on any given day? Mostly no: a human analyst will have an explanation, but if the stock had gone in the other direction, the analyst would have given a different reason. Which means these post-hoc "explanations" are no more valuable than a complex machine-learned model saying, "I predict it will go up because this non-linear function falls on the right side of the decision boundary."

Also, if we really understood hand-coded systems, like operating systems, then they would have no bugs. But they do have bugs, which suggests that we don't fully understand these hand-coded systems—and the main reason we don't understand is that the domain is complicated and we can't envision all the interactions.

Peter Norvig is director of research at Google.

More Spending Won't Solve Anything

Jason Pontin's op-ed calling for more government R&D ("Make America Great Again," May/June 2017) piqued a nerve. I was at a conference at MIT last month in a room full of scientists calling for the same thing, and I had to tell them it's not that easy. I'm a fan of science. I have seven degrees from Stanford and MIT, mostly in STEM, but I'm also a fan of evidence. It's not clear the evidence supports more R&D spending.

As lead energy economist for the White House Council of Economic Advisers in 2006, I was given carte blanche to design a new U.S. energy policy. I was encouraged to spend billions. However, after consulting with the top experts around the world, I realized that there were better ways to spend government money.

Numbers can be bloodless. Eric Schmidt once said at a talk at Stanford that the hardest thing for a CEO is layoffs, because he knew that for every thousand people he laid off, a certain number would commit suicide. So you'd better be sure that those layoffs were worth it. Every government economist holds a similar number in the back of his mind—$8 million. This is the current statistical value of a life used by the federal government.

You are calling for government R&D spending to increase to its peak levels. That's an additional $200 billion, or 27,000 lives per year. We'd better be sure that spending is worth it.

Yes, government R&D as a percentage of GDP has fallen, but it is still more than it was in the past. Adjusting for inflation and using your numbers, we spend $126 billion today compared with $76 billion in the '60s. People regularly call for an "Apollo mission" for climate change, but we already spend more on renewable-energy research than we ever did for the Apollo mission.

People regularly call for an "Apollo mission" for climate change, but we already spend more on renewable-energy research than we ever did for the Apollo mission.

The trouble with pointing to the high returns for existing R&D is that it elides a problem economists spend a lot of time talking about in our intro classes. There is a difference between average returns and marginal returns. When advocates cite studies on the high average returns to government R&D funding, they miss out on this important distinction.

Follow Us Twitter @techreview
Facebook @technologyreview
Instagram @technologyreview

Back when I was an MIT tour guide, the annual research budget for MIT was around $1 billion. Therefore, to spend an extra $200 billion would mean building an extra 200 MITs. That would be amazing, but how would you staff them? The best people in the world are already at MIT. You could start by hiring all the scientists at second-tier schools to fill in your new MITs. Pretty soon those schools would be emptied out and they would start pulling in people from third-tier schools. At the end of the day, all you've done is move people around.

This is one example of the law of diminishing returns. We spend so much on R&D today that spending more is unlikely to do much good.

This is the problem we faced at the White House back in 2006. The most promising areas of research were already well funded. The areas that remained were lackluster. The problem wasn't so much in producing research but in getting those ideas to market.

So the policy process led us to new regulations that would help solve market failures: new regulations for fuel-efficient cars, new regulations for biofuels, loan guarantees for renewable-energy ventures, and yes, some funding for research, which resulted in ARPA-E. Yes, more research funding can be worthwhile, but primarily at the margins, not as the centerpiece of new legislation.

Altogether, these policies became the Energy Security and Independence Act of 2007, which formed the cornerstone of Obama's lauded climate policy.

I do think R&D is worthwhile. I just don't believe in blind faith. We should keep seeking new opportunities to do more, but we shouldn't assume money will solve all our problems.

Ben Ho is an associate professor of economics at Vassar College and Columbia University.

Companies Can't Do It All

In "Make America Great Again," Jason Pontin is right to decry the Heritage Foundation's position that government funding for research is pointless and that private companies should fund such research instead. The idea that private companies can fill the role of creating innovation—whether in their own labs or by funding university research—isn't supported by the facts.

Strategy& (a PriceWaterhouseCoopers unit) studied the 1,000 most innovative companies for more than a decade and found no relationship between financial performance and R&D spending. The problem is that any company large enough to fund research is instead pressured to drastically cut spending and show a return on investment.

Christopher Voigt is the chief technologist at Unisys.

The Evolution of Commenting

Ten years ago, websites had a responsibility to allow people to share their opinions freely in their comments sections. But technology has moved on. Social media is now dominant, and publications no longer need to offer a free and open space for commentary (as Jason Pontin described in "Me and My Troll," May/June 2017). Instead, publishers need to be more strategic about how to utilize community conversation and engagement as a vital part of their journalism. It's a chance to reset who is interacting with them and how. To do this, publishers need to define what success looks like, and engage staff and resources toward that goal. They need to experiment.

Our research at the Coral Project has shown that when a site engages with its audience in the right way, it improves the journalism, creates new community connections, and generates new revenue streams.

Andrew Losowsky is the project lead at the Coral Project at Mozilla.

Twitter Votes

Does Facebook wield too much power in our lives?

46% Yes, it's too powerful

36% I'm not on Facebook

18% No, nothing wrong with it

Would you trust data from a smartphone app to help predict and prevent addiction?

46% Yes, phones are powerful

28% No, it's just an app

26% Only in mild cases

As climate dangers grow, is it time to move ahead with geoengineering field experiments?

48% Yes, we're out of options

37% Only if it's small-scale

15% No, it's too risky

Views



Mary Mossey



Ilan Gur



Bruce Y. Lee

CONNECTIVITY

Virtual Trolls

Social VR has incredible potential. Anonymity has a chance to get in the way of that.

One powerful thing about virtual reality is the way it lets you stand "face to face" with other people—sometimes people from the other side of the world who you'd never otherwise meet. People in social VR settings really tend to connect. I've seen people meet and become friends in VR and then continue the friendship in real life. I've seen people meet their significant others in virtual reality.

But this only works when people appear as their real selves. And one huge problem with social interactions in virtual reality is the degree of anonymity we're finding there. It can alienate people who would otherwise be interested in giving social VR a try.

People using anonymity or maintaining a little more privacy isn't always bad—a shy person might feel able to shake off awkwardness and adopt a more easygoing personality, for example. But what we see too often among the anonymous is bad behavior. With no real name or face it becomes easy to treat people in a critical or abusive way.

We are in the early stages of VR, much like the Internet in the 1990s. There are no rules and no law of the land. We are all learning what is okay and what isn't. And as in those early Internet days, we're learning that anonymity causes a degree of chaos and undesirable behaviors. It removes accountability. Anonymity lets people bully others without repercussions.

There's no simple solution to the problem. Improvements in the areas of friend connections, avatars, and interaction design have helped, but haven't solved it. App creators can help by creating clear standards, including repercussions for bad behavior. Benefits such as customized avatars, extra capabilities, and new features might help coax people to come out of the shadows and engage more authentically.

Whatever we try, it will be worth the effort. Social VR lets people connect and communicate in a way that's much more natural, more effective, and richer in context than current methods of communicating over a long distance. I'd like people to think about VR as a place you can go to express yourself authentically and connect with others. Within that space there's fantastic potential for education, business, science, art, and entertainment. Anonymity kills that potential.

Mary Mossey is a product manager at AltspaceVR, which uses virtual reality to create a new communications platform.

ENTREPRENEURSHIP

The Innovator Gap

The world's most potent technologists are stranded in today's innovation ecosystem.

An innovator I know—let's call him Tom—had trained for a decade at the world's top research institution to become an expert in materials science and engineering. He'd developed a new manufacturing technique for semiconductors with the potential to enable next-generation power conversion devices for lighting, EVs, renewable power, and defense applications. After demonstrating promising results and gathering enthusiastic feedback from industry experts, he was ready to take his technology to market. But how?

I met Tom a year and a half ago, when he applied to an entrepreneurial research program I run called Cyclotron Road. Each year, we receive over 100 applications from the U.S. and beyond for our program, which supports scientists as they work to develop a product based on research in

ANDY FRIEDMAN





PROFESSIONAL EDUCATION

EXPAND YOUR KNOWLEDGE

WITH MIT PROFESSIONAL EDUCATION PROGRAMS

SHORT PROGRAMS

Come to MIT for a Week

Technology professionals, practitioners, and managers across the world come to the MIT campus every summer to learn from MIT faculty who are leaders in their fields ranging from biotechnology and manufacturing to innovation and engineering leadership. Register for a 2-5 day intensive course relevant to your needs and network with your peers from around the globe while gaining critical knowledge to help advance your expertise and impact your organization's success.

Learn more ***professional.mit.edu/short-tr***

Professional Certificate Program in Innovation and Technology

Focused on building competitive advantage through innovation and technology, this professional certificate helps participants acquire timely strategies for developing, managing, and implementing innovative techniques while gaining knowledge of the latest technologies and technical approaches available to industry.

Learn more ***professional.mit.edu/innovation-tr***

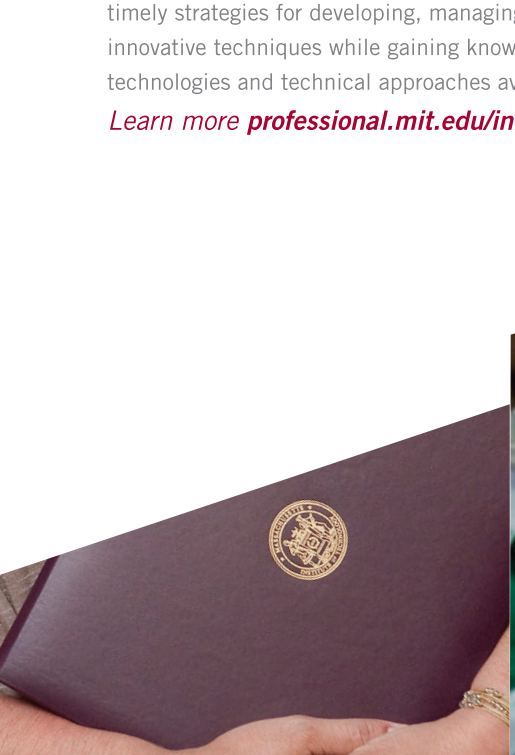
DIGITAL PROGRAMS

Bring MIT to You

Digital Programs allow busy professionals globally to take courses taught by MIT faculty, from any location, and at their own convenience. These online programs focus on important industry topics relevant such as data science, cybersecurity, internet of things, and systems engineering to a variety of professionals around the world.

Learn more ***professional.mit.edu/digital-tr***

Learn more about MIT Professional Education and programs. Visit **professional.mit.edu/techreview2017** or email **professionaleducation@mit.edu**





fields like materials science, chemistry, and physics—what we call "hard technology."

Tom's story is typical. He'd decided to pursue a startup to develop his idea but quickly discovered that the technology was too speculative for private investment. If Tom were a software innovator, he could build a prototype and validate his product in months with little more than a laptop. But to innovate in semiconductor technology, he would need the tools of cutting-edge research—fume hoods, spectrometers, electron microscopes. He would also need time. Turning a semiconductor breakthrough into a market-ready product can take years. Unfortunately, demand by qualified applicants exceeded our capacity, and we had to turn him away.

In February, he sent me an e-mail. After dipping into his personal bank account to get his effort off the ground, he was shifting gears. He'd decided to enter the world of finance and asked if I could introduce him to one of my contacts at a hedge fund. I was saddened, but not surprised.

Today's research ecosystem isn't built to support innovators like Tom. Their projects are too applied for academic labs, which focus narrowly on new scientific discoveries. Meanwhile, private industry can't justify investment in expensive research that doesn't yet have clear commercial potential. Faced with this chasm, even the best innovators struggle. Many are bound to say "why bother?" Can we blame them?

So how do we keep the Toms of this world in the game? For starters, we need to do a better job catalyzing the creation of science-based startups and supporting cutting-edge research within them. Free from the institutional pressures of academic publication or corporate quarterly earnings, startups can serve as powerful vehicles to bring hard technologies to market. Such firms often spend years performing research to develop a first product. Yet despite driving some of the biggest technology disruptions in recent history, small businesses receive only a small percentage of all government research funding.

Society will face enormous challenges in the coming century in areas such as energy, water, food, and health. We need to find a better way to support our best scientists and engineers in their research to turn new discoveries into commercially viable products. Solving these big challenges will depend on it.

Ilan Gur is the founding director of Cyclotron Road.

ROBOTICS

Drones to the Rescue

It's expensive to use drones to deliver meds. Not using them exacts a heavier cost.

If we could send fleets of drones flying medical supplies to people in some of the poorest and most remote parts of the world, would it be worth it? Who's going to pay for the drones, or the equipment needed to house and recharge them? Who'll pay for the personnel to operate them, and the training required?

Last year I was part of a team that tested product-delivering drones in Mozambique with the help of a computational model of the country's vaccine supply chain. The supply chain is the series of refrigerators, freezers, vehicles, and personnel involved in getting vaccines from their origin to the people who need them. The Republic of Mozambique, in southeast Africa, has over 24 million people and a GDP of around $35.3 billion—which turns out to be $419 per capita. Mozambique's 801,537 square kilometers include rugged highlands and mountainous regions, long stretches of plateaus, large rivers, and woodlands filled with wild animals. Heavy rainfalls and cyclones can pummel parts of Mozambique. These conditions can make it dangerous to get vaccines to people.

The "last mile" of the vaccine supply chain—from the district storage facility to the clinics, schools, and homes where vaccines are given—can be tough. Often people on motorbikes, rafts, bicycles, animals, or foot carry vaccines to their final destination. Vaccines are precious because they cost money to develop, can be fragile (heat exposure can make them go bad), and save lives. Sometimes a person has to risk life and limb to carry just a few vaccines to people living far away from any town.

Enter drones. They aren't just for military operations or taking pictures. Working with global health organizations, companies such as Zipline and Matternet are exploring the use of drones to deliver meds.

Our team at the Johns Hopkins Bloomberg School of Public Health and the Pittsburgh Supercomputing Center at Carnegie Mellon University employed software called HERMES that was specially designed to analyze supply chains. Using it, we developed and ran a series of scenarios, including different challenges a drone might face—inclement weather, wild animals, and even people shooting them down out of fear that they were part of a military operation.

The results? Under a wide range of conditions, drones can provide between 20 and 50 percent cost savings over traditional land-based transport. Of course, the savings depend on the size of the population served, the road conditions, the speed of the land-based transport, the reliability of the drones, and the willingness of people to accept them. But it shows that new technology doesn't have to be expensive. By improving service and freeing up valuable resources (like the time of workers who'd normally be carrying vaccines long distances), drones can save both money and lives.

Bruce Y. Lee is an associate professor of international health at the Johns Hopkins Bloomberg School of Public Health.


My Impossible:
To take on my workplace bully, I started a business to take on all workplace bullies.
Catherine Mattice, Civility Partners
Hiscox Customer
I'mpossible
Hiscox business insurance.
Helping the courageous overcome the impossible.
hiscox.com/impossible | 877-505-2048
HISCOX
encourage courage®

Don't miss these must-read stories. Exclusively at
technologyreview.com/mustreadsjuly

must reads

- A Single Self-Driving Car Can Stop a Traffic Jam
- MIT Wants to Jumpstart Reactor Research
- Engineering the Perfect Astronaut
- Could This Apple-Picking Robot Steal Farm Jobs?
- Elon Musk Promises Human-to-Human Telepathy. Don't Believe It.

MIT TECHNOLOGY REVIEW
VOL. 120 | NO. 4

Upfront



MOLLY MATALON

Reinventing Rice for a World Transformed by Climate Change

UC Davis plant geneticist Pamela Ronald wants to create rice varieties that can survive in harsher conditions, including more frequent droughts.

Pamela Ronald stands in front of two rows of rice plants, sprouting from black plastic pots, in a stifling greenhouse on the edge of the University of California, Davis, campus. Researchers in Ronald's plant genetics lab starved the plants of water for more than a week. The ones on the right, the control in the ongoing experiment, are yellowing and collapsing. The

MIT Technology Review

EmTech

Discover the emerging technologies that will change our world.

This is where tech, business, and culture converge...where you gain access to the most innovative people and companies from around the globe.

Register Today:

technologyreview.com/emtech17

November 6-9, 2017

MIT Media Lab
Cambridge, MA

THE HARTFORD

LEMELSON·MIT
Celebrating invention, inspiring youth

Fitzpatrick | we are IP.

Meet the experts and join the conversation.

- The Future of Work
- Can AI Get Smarter?
- Adapting to the Reality of Climate Change
- Social Media and Society
- Next-Generation Brain Interfaces
- Meet the 2017 Innovators Under 35

Featured Speakers Include:



Kris Hammond
Chief Scientist and Cofounder, Narrative Science



Joi Ito
Director, MIT Media Lab



David Keith
Professor, Harvard University
Founder, Carbon Engineering

Upfront

adjacent plants, equipped with an added gene, are thick, tall, and green.

The hope is that these or similar genetic alterations could help rice and other crops survive devastating droughts, preventing food shortages in some of the poorest parts of the world. Ronald, a trim scientist with short brown hair, smiles as she looks down at the early results.

She has spent the last three decades working to make rice, a food staple for more than half the world's population, more resistant to environmental stress. She was a central player in one of the greatest recent success stories in plant genetics, isolating a gene that allows rice to survive extended periods of flooding. It's a huge challenge in low-lying parts of Asia, wiping out around four million tons of rice each year in India and Bangladesh alone. A decade after her lab's discovery, more than five million farmers grow rice varieties engineered with the so-called *Sub1* gene, covering more than two million hectares across Asia.

The latest research could be even more significant as climate change ratchets up the frequency and intensity of droughts across large swaths of Earth, threatening the food security and stability of entire nations. The number of extreme droughts could double by the end of the century, devastating fields and farmers across South Asia and sub-Saharan Africa.

Ronald's work provides a powerful statement about the potential of modern genetic tools to preserve livelihoods and lives, offering a counternarrative to the widespread fears and distortions surrounding genetically modified crops. "This focus on genes in our food is a distraction from the really, really important issues," she says. "How can we reduce the use of toxic inputs? How can we feed the poor and malnourished? How can we be sure that farmers have access to seeds, and that consumers can afford the food?"

Ronald grew up San Mateo, California. Her mother was a dedicated gardener, cook, and environmentalist. Her father was a businessman who fled Nazi Germany as a child. Years after arriving in California, they built a 500-square-foot cabin in south Lake Tahoe, where the family spent summer vacations. One hot day when she was around 15, Ronald and her brothers hiked a steep path into the High Sierra. At the saddle, they happened upon a couple hovered over a book. They were a pair of professional botanists who were cataloguing flowers. She had developed an affection for plants from the time she spent with her mother in the garden and kitchen, but this was the first time she realized you could make a living working with them. In the late 1980s, during her PhD program at UC Berkeley, Ronald started working with peppers and tomatoes. But as she began her postdoctoral work, she decided to shift her focus to rice, realizing that even small advances in stress tolerance for such a critical crop could help a lot of people. Tomatoes and peppers are "important for salad," she says, "but I wanted to work on supper. I wanted to work on a staple food crop."

The number of extreme droughts could double by the end of the century.

Ronald arrived at UC Davis as an assistant professor in 1992. Her work on flood-tolerant rice started in the mid-1990s, as a U.S. Department of Agriculture–funded collaboration with colleagues. Over the course of a decade, the team pinpointed and isolated the *Sub1* gene in an ancient but unpopular Indian rice variety that is able to survive even after being submerged underwater for more than two weeks. Since then, the Philippines-based International Rice Research Institute, backed by more than $70 million in funding from the Bill & Melinda Gates Foundation, has bred that gene into 10 popular Asian rice varieties. In turn, the nonprofit put the seeds into the hands of farmers in India, Bangladesh, Indonesia, Nepal, and other nations.

Rice is a tough crop to grow, requiring a lot of work and a lot of water. Too much water all at once kills it, but so does too little. It takes just a week without rain to significantly decrease yields in hilly rice-growing areas. The challenges of rice production are only bound to get worse in many areas as climate change raises temperatures, reduces rainfall in certain places, and increases flooding or sea-level rise in others. Under a climate scenario that assumes high greenhouse-gas emissions, rice yields would be nearly 15 percent lower than otherwise expected at midcentury, and prices would be 30 percent higher, according to a 2015 report in *Environmental Research Letters*.

Shifting farming practices and the fertilizing effect of increased carbon dioxide could offset some of these climate impacts. But it's going to become much harder and more expensive to maintain yields in many areas, and rich nations will be in a far better position than poor ones to make the necessary changes, says Keith Wiebe, a senior research fellow at the International Food Policy Research Institute.

The work at Ronald's lab on drought-tolerant rice varieties is in an early phase. She declines to discuss details, including the basic approach, until she and her colleagues have verified the initial results in additional experiments and published their findings. Other scientists around the world are also racing to develop drought-resistant crops and have already achieved some advances, including sprays, hybrids, and genetic alterations that help crops get

by with less moisture, perhaps by switching into water-preserving modes at earlier signs of trouble. But greater advances will be required to confront the growing challenges ahead, and drought tolerance is a tricky problem. The trait generally involves various genes and cellular communication pathways. It's crucial that any improvements not come at the expense of yield, taste, and other qualities important to farmers and consumers. And there would seem to be hard limits on how much can ever be achieved, as all plants need water.

On an overcast Saturday in late April, Ronald stood on stage at a brick plaza on the edge of the San Francisco Bay, addressing a sign-wielding crowd gathered for the March for Science. "Science is based on data, not on alternative facts," she said, pausing for applause at the end of most sentences. "Science is not a buffet where people can pick and choose the parts they like, and throw out the rest." But people do, of course. The weakest applause line of her speech before the crowd, gathered largely to protest the Trump administration's denial of climate science, came when she said that science had improved California's fruits, vegetables, and nuts. In other words, when she acknowledged a field that could help address some of the problems arising from a changing climate.

It was typical of Ronald, who's determined to say where she believes the science leads, whoever the audience may be. Genetically modified crops have become incredibly contentious, widely portrayed as reckless attempts to tinker with Mother Nature for the sole benefit of seed conglomerates. But the body of science shows they've been both safe and beneficial, Ronald argues. She publicly sparred with the Union of Concerned Scientists on these issues, suggested that Greenpeace was "misinterpreting data," and criticized Vermont's GMO labeling laws.



Pamela Ronald gets hands-on with a rice plant.

Ronald's research highlights the promise for genetic alterations. *Sub1* rice sidestepped any anti-GMO backlash because, while the tools of modern genetics were necessary to isolate and express the gene, it doesn't carry any non-rice DNA. The trait from one rice variety was added to others through modern breeding methods, accelerated by analyzing the DNA of offspring to avoid false paths.

Ronald notes that every major food crop has been altered by human hands in one way or another. And some of the most important advances in the future, to improve yields, nutrition, or environmental tolerance, may be possible only with genetic engineering and powerful gene-editing technologies such as CRISPR. What should matter to lawmakers, regulators, or critics isn't which implement was pulled from the ever-advancing genetic toolbox, but whether it produced a positive or negative impact on human health or the environment. At this point, we have a four-decade track record of genetic engineering in plants, medicine, and cheese, with no evidence of harm, Ronald says. The danger is that unfounded fears could prevent us from easing real human suffering—for example, if misguided regulations slow down the science, or protests prevent seeds and crops from reaching the farmers and consumers who need them most. For Ronald, the real goal should be sustainability in the broadest sense, applying whatever combination of breeding, organic farming, and genetic technology helps us feed a growing population without exacting a higher environmental cost. "We need to make policy based on evidence, and based on a broader understanding of agriculture," she says. "There are real challenges for farmers, and we need to be united in using all appropriate technologies to tackle these challenges."

—James Temple

MOLLY MATALON

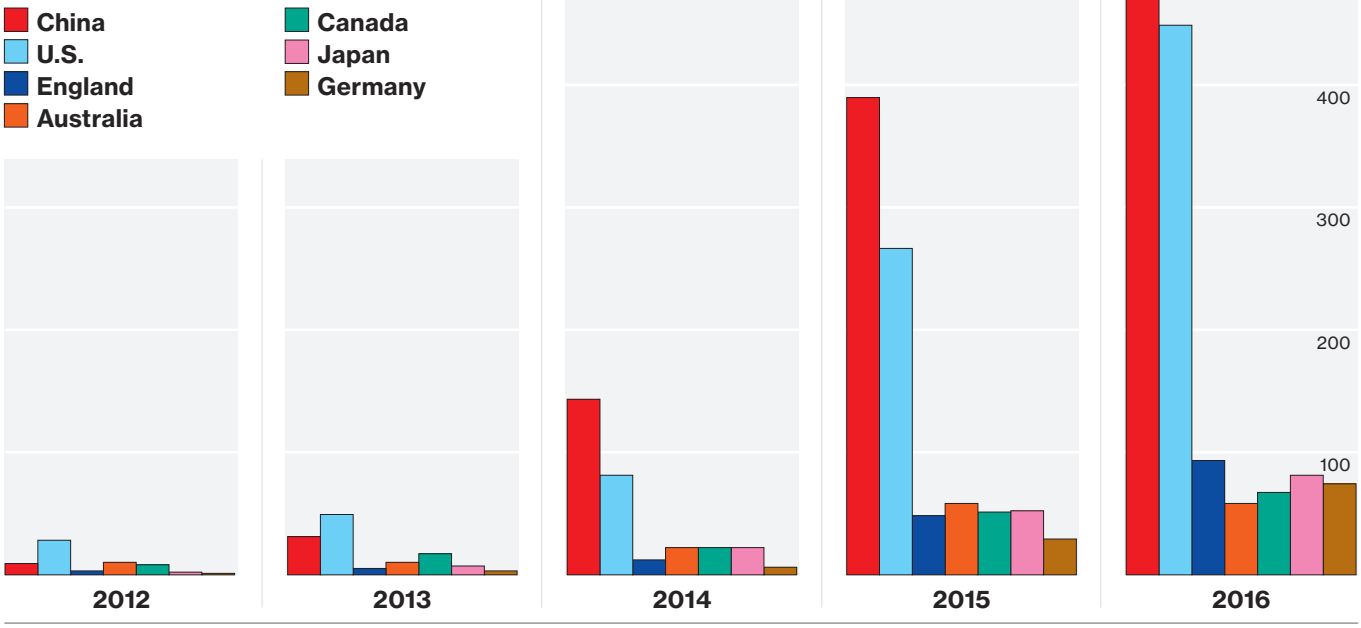
Upfront

Who Is Winning the AI Race?

China and the United States dominate the world of artificial-intelligence research. Microsoft, IBM, and Google are the leading companies.

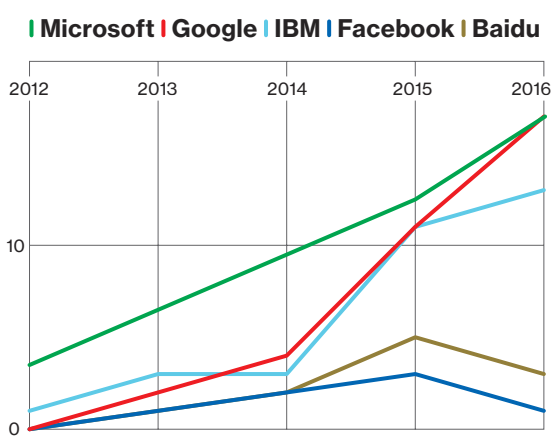
China Learns Quickly

Since 2014 China has published the most research papers per year on deep learning, an advanced form of artificial intelligence.



The Big Three

For years, Microsoft published the most deep-learning research papers, but Google and IBM have gained ground.



The Fight for IP

IBM has dominated U.S. patent activity in AI, but Google and Facebook may be closing the gap. Since the process involves a time lag before applications are published, records from 2014 onward are probably not complete.

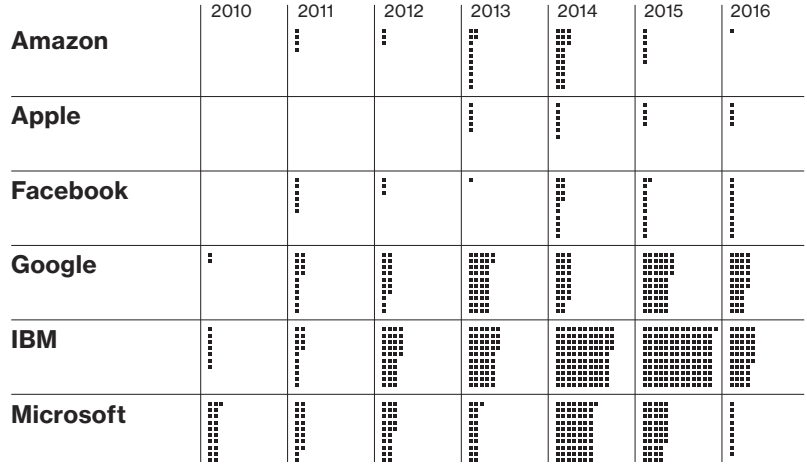
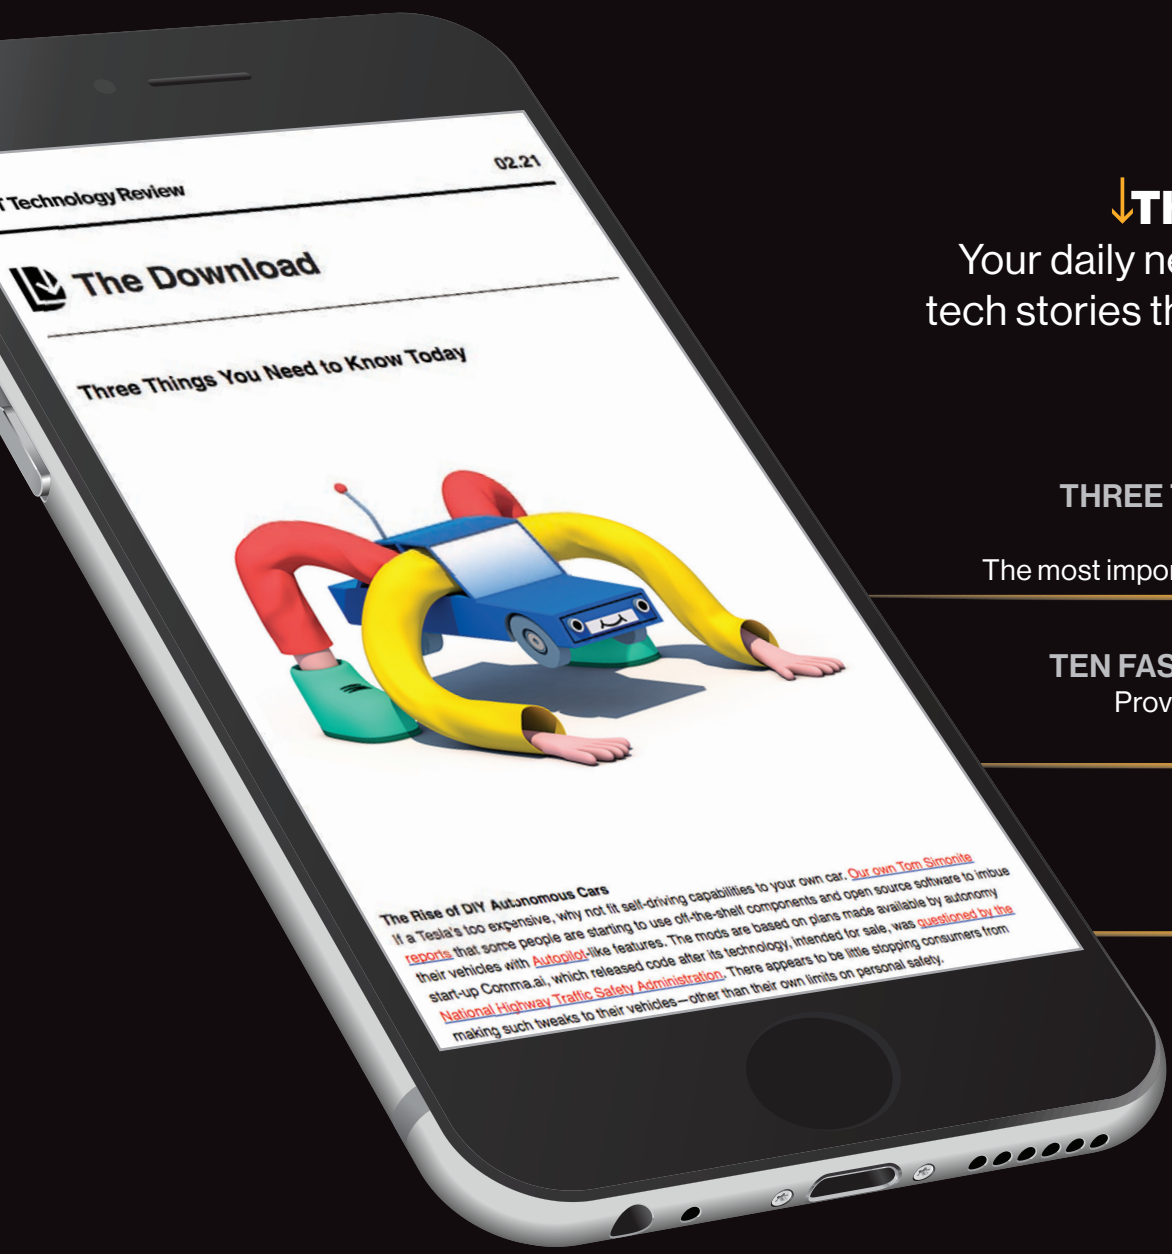


ILLUSTRATION BY LUKE SHUMAN; RESEARCH PUBLICATION DATA FROM THE WEB OF SCIENCE; PATENT DATA FROM CB INSIGHTS

Get MIT Technology Review delivered to your inbox

EVERY DAY.



↓The Download
Your daily newsletter of the tech stories that matter most

THREE THINGS YOU NEED TO KNOW TODAY
The most important stories of the day

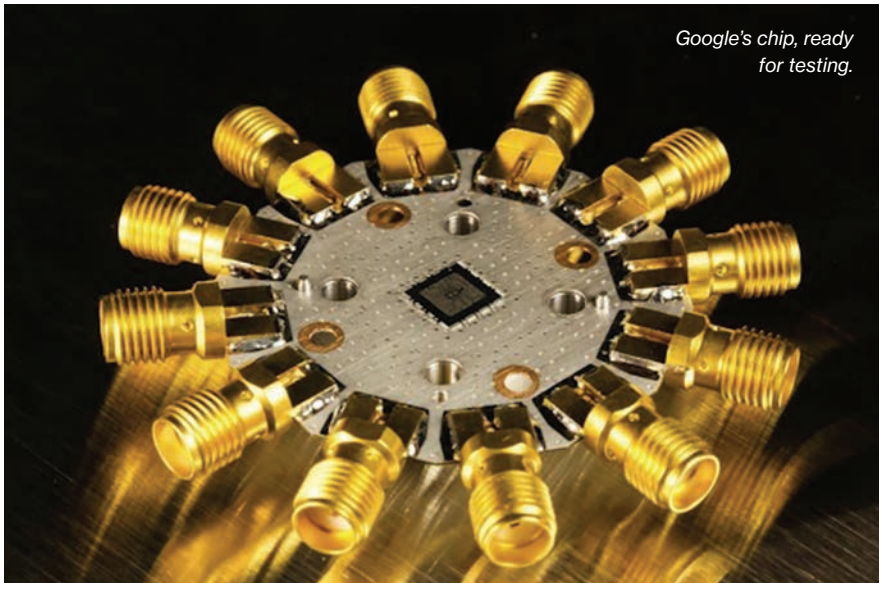
TEN FASCINATING THINGS
Provocative items from the world of technology

AND MUCH MORE

Sign up now:
technologyreview.com/everyday

MIT Technology Review

Upfront



Google's chip, ready for testing.

Google's Quantum Leap

The search giant plans to build a chip that will achieve "quantum supremacy" in computing by year's end.

John Martinis has given himself just a few months to reach a milestone in the history of computing. He's leader of the Google research group working on building astonishingly powerful computer chips that manipulate data using the quirks of quantum physics. By the end of this year, Martinis says, his team will build a device that achieves "quantum supremacy," meaning it can perform a particular calculation that's beyond the reach of any conventional computer. Proof will come from a kind of drag race between Google's chip and one of the world's largest supercomputers.

One reason for his confidence is that Google's roughly 25-strong group has made a new quantum chip that tests out key design features needed to make a device for that head-to-head contest. Quantum chips represent digital bits of data using qubits, devices that can shortcut through some tough calculations by exploiting the counterintuitive physics of quantum mechanics. Researchers have so far demonstrated quantum computing with only small groups of qubits, though. Google has released results from a chip that has nine qubits arranged in a line, but Martinis says he'll need a grid of 49 qubits for his quantum supremacy experiment. Google's latest chip has only six qubits, but they are arranged in a two-by-three configuration that Martinis says shows the company's technology still works when qubits are nestled side by side, as they will be in larger devices.

The race to develop quantum processors includes Intel, Microsoft, and IBM.

The six-qubit chip is also a test of a manufacturing method in which the qubits and the conventional wiring that controls them are made on separate chips later "bump bonded" together. That approach, a major focus of Google's team since it was established just over two years ago, is intended to streamline the extra control lines needed in a larger chip, which can interfere with how qubits function.

"That process is all working," says Martinis. "Now we're ready to kind of move fast." Designs for devices with 30 to 50 qubits are already in progress, he says. Martinis joined Google in late 2014 from the University of California, Santa Barbara, where he remains a professor today. His team is one of several industrial research groups recently formed or expanded thanks to growing signs that the technology behind quantum computing is becoming more tractable. The race to develop quantum processors includes Intel, Microsoft, IBM, and even startups. Simon Gustavsson, a principal investigator in a quantum computing research group at MIT, says that Google is one of the leaders. "It's pretty comparable between Google and IBM," he says.

Pulling off its quantum supremacy experiment this year would underline the search company's competitiveness, although quantum processors would need to be much larger than 50 qubits to be capable of useful work. "It'll be an academic milestone," says Chris Monroe, a professor at the University of Maryland and cofounder of the quantum computing startup IonQ. "Afterward you still have to figure out how to make it more scalable and programmable."

Martinis agrees that much will remain to be done, but he argues that the experiment could become a benchmark for anyone claiming to have a working quantum computer. *—Tom Simonite*

COURTESY OF GOOGLE

Edible CRISPR Could Precisely Target Dangerous Germs

Researchers are in the early stages of developing a probiotic to make disease-causing bacteria self-destruct.

As resistance to antibiotics grows in the U.S., researchers are looking for new ways to fight germs like *Clostridium difficile*, a bacterium that can cause fatal infections, often in hospitals and nursing homes. One way to do that could be with a "CRISPR pill" that instructs harmful bacteria to self-destruct.

CRISPR is the powerful gene-editing technology that's already being explored as a way to precisely edit human genes to cure diseases. Because of its versatility, it's also being studied for a huge range of other uses. One possibility is to turn it into ultra-precise antimicrobial treatments to "specifically kill your bacteria of choice," says Jan-Peter Van Pijkeren, a food scientist at the University of Wisconsin, Madison. While not a household name, *Clostridium difficile* tops the U.S. Centers for Disease Control and Prevention's list of urgent drug-resistant threats. A 2015 study by the agency found that the bacteria caused nearly half a million infections in Americans, leading to 15,000 deaths.



Clostridium difficile *bacteria in a petri dish, glowing under long-wave ultraviolet light.*

CRISPR was actually discovered in bacteria. In fact, the system is an immune defense bacteria use to fend off invading viruses called bacteriophages. Bacteria store memories of viral DNA in their own genomes as "clustered regularly interspaced short palindromic repeats." They use this memory, plus a DNA-slicing enzyme known as Cas, to recognize and chop up the genes of invading bacteriophages. Van Pijkeren's idea is to use a bacteriophage to send a false message to *C. difficile*, one that causes the bacteria to make lethal cuts to its own DNA. To do it, his lab is developing bacteriophages capable of carrying a customized CRISPR message. To get the viruses into a person, Van Pijkeren plans to add them to a cocktail of innocuous bacteria, or probiotic, that a person could swallow as a pill or a liquid. As it passes through a person's intestinal tract, the bacteriophage would infect any nearby *C. difficile*, causing the bacteria to hack up their own DNA.

Van Pijkeren says the probiotic is still in early stages of development. However, researchers have previously shown that using bacteriophages to trigger CRISPR can efficiently kill skin bacteria and might also help combat a diarrheal infection common in the developing world.

The appeal of using CRISPR is that such drugs would be very specific—theoretically, they would kill a single species of germ while leaving beneficial bacteria intact. Broad-spectrum antibiotics, by contrast, kill off large swaths of both good and bad bacteria. In fact, the overuse and abuse of conventional antibiotics is what leads to resistance in the first place.

—Emily Mullin

TO MARKET

The Light Phone

A Dumb Phone

COMPANY:
Light

PRICE:
$150

AVAILABILITY:
Now



COURTESY OF LIGHT; COURTESY OF MIT

Dependence on smartphones can feel like a shackle, and sometimes you just want to get away from that constant connectivity. The Light Phone promises just that. When someone dials your regular smartphone number, the call is forwarded to your Light Phone via an online server. Calls you place look as if they're coming from your usual number, too. The credit-card-size Light Phone has its own SIM card, so it doesn't need to be anywhere near your smartphone to function. It's barely recognizable as anything more than a thin slab of plastic until you press the power button, at which point it lights up with a number pad and single-line display. It is a phone in the simplest sense. *—Rachel Metz*

Upfront

How a Boy's Lazarus-like Revival Points to a New Generation of Drugs

Drugs made from RNA may be the next great class of medicine.



Cameron Harding is one of the first children to survive spinal muscular atrophy.

LEIGH WEBBER

Shortly after Cameron Harding's one-month check-up, his mother, Alison, saw that her newborn seemed to stop moving. She'd unwrap him from a swaddle and his arms would flop to one side. He wouldn't kick his legs or turn his head. The diagnosis: spinal muscular atrophy. The inherited illness, which destroys the motor neurons that control movement, often kills children before they turn two. In Cameron's case, it seemed he'd probably never even have a birthday.

But when he was seven weeks old, Cameron's parents enrolled him in a clinical trial for an experimental drug. In videos shot two months later, he could move his head and reach for a toy. No child with his condition had ever made such a recovery before. The drug Cameron received, Spinraza, was developed by Ionis Pharmaceuticals and Biogen and was approved in the U.S. late last year. It may become the first blockbuster in a novel category of drugs called RNA therapeutics, after the genetic messenger molecule from which they are constructed.

Drugs made of RNA have been in development for more than 20 years. But Cameron's remarkable trajectory is a sign they could be ready to join familiar pharmacy chemicals and biotech proteins as the next great class of drug types, proponents say. "Right now, RNA therapeutics—that's the future of medicine," says Steven Dowdy of the University of California, San Diego, School of Medicine. Dowdy admits to being a partisan—his lab works on RNA—but says advances in chemistry have finally made this type of drug feasible. "We still have a lot of work to do, but it's going to be huge," he says.

Some investors think so, too. One private company, Moderna Therapeutics, has reportedly managed to raise nearly $2 billion to develop what it calls a "platform" for inventing new RNA treatments. Overall, more than 150 clinical trials under way are using RNA to treat cancer, infection, hormone problems, and neurologic diseases, including Huntington's disease.

"This has given us a glimpse into a future that's really quite profound."

Inside our cells, the genes in our DNA code are translated into copies made of RNA. These copies, in turn, float into the cell body, where they serve as the information template from which proteins are manufactured. Most drugstore pills act directly on proteins. Gene therapy, meanwhile, seeks to replace DNA instructions with new ones.

What drugs like Spinraza seek to do instead is use RNA to block, modify, or add to the existing RNA messages in a cell. The most serious challenge to this approach, says Dowdy, is that cells evolved billions of years ago to keep RNA from the environment out. Overcoming this delivery problem has cost company chemists many years of effort.

While the science of RNA has unfolded since the 1960s to a drumbeat of Nobel Prizes, the advance into real medicine hasn't been quick. RNA drugs have been beset by serious side effects, and several have been pulled from human tests over safety concerns. In October, the biotech company Alnylam had to stop a large study of one drug, Revusiran, after unexplained deaths in volunteers being treated for a hereditary metabolic disease, ATTR amyloidosis. Alnylam, which works with a technique for blocking RNA mes-



sages, has spent about $1.5 billion since being founded in 2002 but is still working toward its first product. In a January paper, it showed that one of its drugs could lower cholesterol levels for six months with one shot.

Rachel Meyers, previously a senior executive at Alnylam, says one problem facing RNA drugs is that companies have too often tried to use them to treat diseases where other options exist. Instead, she believes, RNA therapeutics are more likely to find success in situations, like Cameron's, that can't easily be treated in any other way. Spinraza works through a mechanism called "antisense." The drug is a chemically enhanced strand of RNA that, by matching up with a mirror copy present in Cameron's cells, allows his body to correctly assemble a protein his nerve cells need.

Last year, two drugs that work along these lines were approved in the U.S. The other, Exondys 51, developed by Sarepta Therapeutics, was approved to treat muscular dystrophy after becoming the focus of impassioned lobbying by parents of affected boys, who prevailed on the U.S. Food and Drug Administration to allow it on the market despite limited evidence of its benefits.

There are no such questions about the drug Cameron receives. Its effects are "Lazarus-like," according to one expert, but it is not simple and is not a permanent cure. It's delivered through a spinal tap once every four months, and Cameron will need it for the rest of his life. The cost of the injections: $375,000 a year. "This has given us a glimpse into a future that's really quite profound, at least in serious disorders of the central nervous system," says Michael Ehlers, executive vice president of Biogen.

What happens next for Cameron is not clear. Because Spinraza just started being tested in children five years ago, doctors aren't sure what will occur as Cameron grows older. C. Frank Bennett, senior vice president of research at Ionis, says the company is now studying the results if the drug is given even earlier, just after birth, and before symptoms start. So far, he says, these kids are meeting their developmental milestones by rolling over and crawling on time. *—Karen Weintraub*

QUOTED

"Software is eating the world, but AI is going to eat software."

— Nvidia CEO Jensen Huang on how the machine-learning revolution is just beginning.

"We have no interest in creating more survivors with impaired quality of life."

— Alan Flake, a surgeon at Children's Hospital of Philadelphia, where researchers kept eight premature baby lambs alive in artificial wombs for weeks.

"I would block its attempts to patrol my facilities, to prevent it from monitoring my assets."

— 24-year-old Byun Hyun Woo, a champion Starcraft player, who does not believe artificial intelligence will master the game in his lifetime.

BY THE NUMBERS

500

Number of new cars that China's BYD, the largest electric-vehicle manufacturer in the world, is making per day, down from 900 before the government cut subsidies for EVs.

29

Number of hours a home can be powered by a single revolution of the world's largest wind turbines, 195 meters tall, which are operating off the northwest coast of England.

10,000

Factor by which the cost of DNA-based data storage needs to fall before it can become commercially available, according to Microsoft, which plans to be ready to offer it when it does.

$150,000

Annual cost of the immunotherapy Keytruda, the first drug approved to treat cancer on the basis of its genetic makeup rather than its location in the body.

Upfront



Real or Fake? AI Is Making It Very Hard to Know

Thanks to machine learning, it's becoming easy to generate realistic video, and to impersonate someone.

News headlines might not be the only things that are fake in the future. Powerful machine-learning techniques are making it increasingly easy to manipulate or generate realistic video and audio, and to impersonate anyone you want with amazing accuracy.

A smartphone app called FaceApp, released recently by a company based in Russia, can automatically modify someone's face to add a smile, add or subtract years, or swap genders. The app can also apply "beautifying" effects that include smoothing out wrinkles and, more controversially, lightening the skin. And a company called Lyrebird, which was spun out of the University of Montreal, has demonstrated technology that it says can be used to impersonate someone's voice. The company demonstrated it with clips in which Barack Obama, Donald Trump, and Hillary Clinton apparently endorsed the technology. These are just two examples of how the most powerful AI algorithms can be used for generating content rather than simply analyzing data.

Powerful graphics hardware and software, as well as new video-capture technologies, are also driving this trend. Last year researchers at Stanford University demonstrated a face-swapping program called Face2Face. It manipulates video footage so that a person's facial expressions match those of someone being tracked using a depth-sensing camera. The result is often eerily realistic.

Both FaceApp and Lyrebird use deep generative convolutional networks to enable such tricks. They are applying an emerging technique that lets algorithms go beyond just learning to classify things and generate plausible data of their own by using very large, or deep, neural networks. Such networks are normally fed training data and tweaked so that they respond in the desired way to new input. For example, they can be trained to recognize faces or objects in images with amazing accuracy. But the same networks can then be made to generate their own data. For instance, such a network can generate images from scratch that look almost like the real thing. In the future, using the same techniques, it may become a lot easier to manipulate video, too.

—Will Knight

TO MARKET

Lighthouse

Home Security Assistant

COMPANY:
Lighthouse

PRICE:
$399

AVAILABILITY:
September 2017



A new smart-home assistant and security monitor can tell the difference between specific adults and spot kids and pets, and send you smartphone alerts about what they're up to. Lighthouse uses two cameras, including a 3-D time-of-flight camera that can see how far away an object is and distinguish objects in the foreground from those in the background. If the device finds something that may be interesting—say, your kids walking into the living room at 11 P.M., or an unknown person in the house while you're out—it will send the information to a remote server, which analyzes the data and works with a Lighthouse app running on your smartphone to figure out what to do with it. *—Rachel Metz*

GWENDAL LE BEC; COURTESY OF LIGHTHOUSE

Self-Driving Cars' Spinning-Laser Problem

Progress toward practical autonomous vehicles requires improvements to the sensors that map a vehicle's environment in 3-D.

Many components go into making a vehicle capable of driving itself, but one is proving to be more crucial and contentious than all the rest: the lidar sensor, a device that maps objects in 3-D by bouncing laser beams off its surroundings. Self-driving vehicles being tested by companies such as Alphabet, Uber, and Toyota rely heavily on lidar to locate themselves on the detailed maps they need to get around, and to identify things like pedestrians and other vehicles. The best sensors can see details of a few centimeters at distances of more than 100 meters.

Most companies in the race to commercialize self-driving cars consider lidar essential (Tesla is a rare exception, relying solely on cameras and radar, which can't map the world in 3-D anywhere near as accurately). But self-driving technology has ramped up so fast that the nascent industry is suffering from a kind of lidar lag. Making and selling lidar sensors was previously a relatively niche business, and the technology doesn't yet seem mature enough to become a standard component in millions of cars.

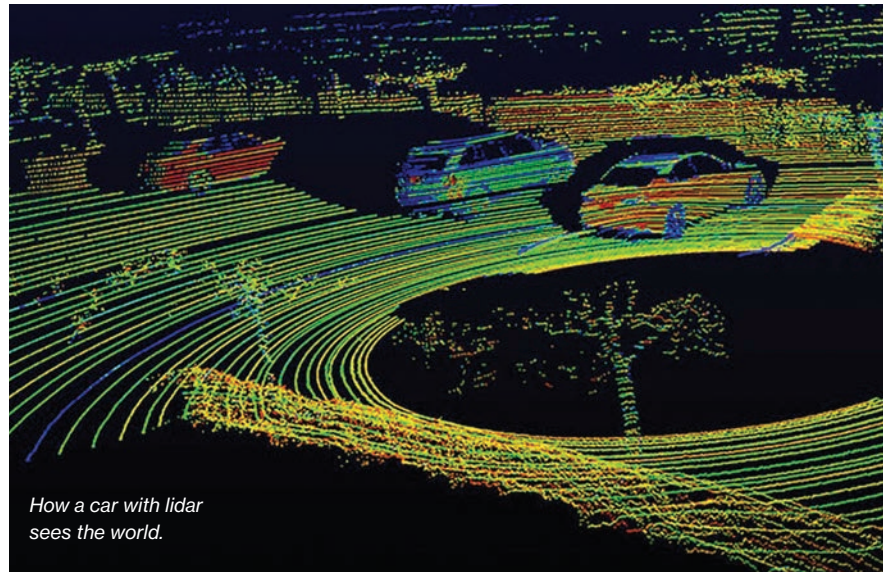
One problem is that lidar sensors are bulky. They are why vehicles being tested by Waymo, Alphabet's self-driving-car unit, are topped by a giant black dome, and Toyota's and Uber's sport spinning gadgets the size of a coffee can. They're also expensive, costing thousands or even tens of thousands of dollars apiece. Most vehicles in testing have multiple lidars onboard, and despite the relatively small numbers on the road, demand has become a problem for lidar manufacturers.

It's a problem that's getting plenty of attention. Better lidar is a core part of Waymo's plan to make self-driving cars a mass market and a profitable proposition. The company has developed three different sensors that look for objects at different ranges. And last year Ford and Baidu, the Chinese search company investing in self-driving cars, jointly invested $150 million in Velodyne, the world's leading lidar supplier. The company is building a new "megafactory" in San Jose that's scheduled to start churning out lidars starting next year.

Still, many in the self-driving-car industry think lidar needs reinventing if it is to become practical enough. Velodyne is one of several companies working on designs that don't use spinning mirrors to direct their laser beams out into the world, as the devices on the road today do. Versions that steer their lasers electronically, described as solid state, should be much cheaper, smaller, and more robust, because they don't have moving parts.

It's a theory yet to be fully tested. Velodyne reported last December that its project had made a "breakthrough" that could make lidars as cheap as $50, and it plans to start producing a solid-state sensor next year. Startup Quanergy, which last year scored $90 million in funding, says it will start producing solid-state lidar sensors at a factory in Massachusetts this year and sell them for $250, but full details are still unclear. Auto-parts suppliers Continental and Valeo are working on similar technology of their own, but they say it will come to market in two or three years.

Automakers including Ford and BMW have said they want to have fleets of autonomous cars operating on roads by 2021. The performance, cost, and looks of those vehicles will be shaped by progress on the lidar sensors so crucial to today's prototypes. *—Tom Simonite*



How a car with lidar sees the world.

COURTESY OF VELODYNE LIDAR

Why Tesla ⇧ Is Worth More Than GM ⇩

A few companies that master digital technologies are capturing huge chunks of the economy. Does this explain the persistence of slow growth?

By James Surowiecki

MIT TECHNOLOGY REVIEW
VOL. 120 | NO. 4

The digital economy has transformed the way we communicate with each other; the way we consume information, products, and services; the way we entertain ourselves. It's revolutionized seemingly non-digital industries—think of how different financial services, for instance, are today from what they were two decades ago—and investors expect it to soon transform others, which is why Tesla Motors is worth more than General Motors despite making a tiny fraction as many cars as GM makes and earning a tiny fraction of the revenue.

This phenomenon explains why the so-called Big Five of the digital economy—Apple, Alphabet, Microsoft, Amazon, and Facebook—have, at various points over the last year, been the five most valuable companies in the world.



So you might say that the digital economy has lived up to the expectations people had for it 20 years ago, in the early days of the Web. In other important ways, however, its consequences have been smaller than you might think. U.S. GDP growth has, by historical standards, been disappointingly slow since the arrival of the Internet. Productivity growth, which many assumed would be invigorated by the impact of digital technologies, has been dismal for much of this century. In the heyday of the Internet boom during the late 1990s, productivity growth accelerated for the first time since the 1970s, and it appeared for a short while as if technological innovation had solved a central problem with the U.S. economy. But the productivity boom ended in the early 2000s, and it never resumed. Some observers initially suspected mismeasurement, arguing that GDP was not capturing the true value of the many free goods the digital economy offers. But there's little doubt that the productivity revolution we hoped digitization would usher in has yet to materialize.

The digital economy also has not transformed the job market as much as one might have expected. To be sure, we now have entirely new categories of workers: fleets of Uber drivers, and the TaskRabbiters who haunt Whole Foods stores in major cities. But Americans aren't job-hopping more than they used to—in fact, by some measures they're changing jobs less than at any point in the past two decades. And digitization has also eliminated swaths of workers—not just because of automation, but also because of things like online shopping, which has put hundreds of thousands of retail workers out of a job. More important, the digital economy has not been the source of a huge number of good, well-paying jobs. In fact, the rise and consolidation of the digital economy has coincided with an extraordinarily weak labor market. Wages for American workers have just recently started to grow at a clip faster than inflation, but for most of this century they've been close to stagnant. (The same is true in most developed countries.) That's not the fault of digitization. But digitization has not been the driver of job and economic growth that many hoped it would be. Information and communication technologies—which include software and IT companies and Internet firms, along with entertainment and publishing—have seen their share of GDP rise just 1 percent since 2000. And while that's almost certainly an understatement—and doesn't capture the impact of digitization on other industries—the number is striking. So, too, is the fact that only a small percentage of private-sector workers work in what you would think of as digital companies.

Still, the most surprising—and potentially troubling—thing about today's digital economy is how remarkably stable it has become. The buzzword that was always associated with digitization was "disruption." The Internet and other digital technologies, it was assumed, would accelerate competitive pressures and make it harder for incumbents to hold onto power. If the old industrial order was characterized by companies that stayed at the top for decades, the digital economy, with its supposedly low barriers to entry and low switching costs, was going to be characterized by constant turnover at the top. Instead, the opposite is true. Today's digital economy, at least on the consumer side, is dominated by the same five giants that have dominated it for at least the past decade and that almost everyone seems to anticipate will dominate it for the foreseeable future (at least if you go by their market capitalizations, which anticipate many more years of enormous profits for all of them). The digital economy is an economy in which platforms are the biggest source of value, and the Big Five's platforms are the most lucrative ever invented. The result is that this economy is governed, in effect, by an oligopoly. The Big Five sometimes compete and sometimes coöperate, but ultimately each has solid control over its core markets.

"Oligopoly" sounds sinister, but this one was not primarily created by overtly anticompetitive or monopolistic behavior. Rather, digital markets are what economists call winner-take-all markets, in which success tends to breed nearly insuperable advantages. The rule that seems to govern today's digital economy, in fact, was well stated in Matthew 13:12: "For whoever has, to him more shall be given, and he will have an abundance." Which is great for those who have, and not so great for everyone trying to compete with them.

Power of numbers

How did we end up with a digital economy dominated by a few big players? The simplest explanation focuses on what are called network effects, whereby a product or service becomes more valuable the more people use it. In the classic example of a network effect, a telephone is worthless if only one person has one, since there's no one to call. If two people have telephones, they now have some value. And if a million people have telephones, the phone network suddenly becomes enormously valuable.

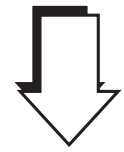
The implication is that the more users a network has, the easier it becomes to add more users. And direct network effects in this sense are important in understanding the success of a company like Facebook. Facebook's single biggest advantage over any would-be competitors, at this point, is simply that it's such an immense network that if you want to connect with people, it's the logical place to start. The same is true of services like Instagram and China's WeChat. For digital companies like Snap and Twitter, which are struggling to become profitable, direct network effects are just about the only value they have at all.

The Big Five also benefit from what are sometimes called indirect network effects, including the fact that sellers want to be where buyers are, and vice versa. Because Google has such an enormous user base, companies want to advertise with it. And that makes Google a natural place to go if you're looking to buy something. Similarly, because Amazon has such a critical mass of customers, it's the natural place for third-party sellers to gravitate. When Amazon made the decision to allow third-party sellers on its site, competing with its own wares, the decision seemed crazy to many at the time. But it positioned the company to benefit from the network effect: having third-party sellers made Amazon more appealing to customers, which in turn made it more appealing to sellers, creating a virtuous cycle for the company.

Beyond the network effects is another, related way that the sheer scale of the Big Five helps them stay on top: through the access they have to enormous amounts of user data. That data, which is far more detailed and granular than anything companies have been able to access in the past, helps these companies improve their products and services, which in turn helps them add more users, which gives them access to more data, and so on. This data flywheel effect did not get as much attention as network effects in the early days of the digital economy, but it's become clear that user data is an enormous competitive advantage for the powerhouses of the digital economy and a key reason why it's difficult to imagine them being toppled anytime soon. By tracking people's clicks, Google is continually improving its search results and ad servicing. Amazon and Netflix and Apple mine their data to improve their recommendation algorithms, making it more likely they'll offer you things you want to buy or watch. This process isn't automatic—you need to have lots of smart data scientists and be willing to invest the resources to ceaselessly upgrade your product. But if you do that, and all of the Big Five do, the rewards can be immense—far greater than what comes merely from the traditional online business model of packaging data and selling it to advertisers.

The ability to gather enormous amounts of data and analyze it efficiently is also at least part of the reason investors think Tesla is more valuable than General Motors. After a traditional car company sells a car to a customer, its relationship with that customer typically is limited (except for maintenance and servicing). Tesla, by contrast, collects terabytes of driving data—including, in some cases, video data—from its customers. That data is then put to use in improving the self-driving features of its cars. According to Adam Jonas, an analyst at Morgan Stanley, Tesla cars are now logging five million miles a day. Since making self-driving cars work depends on machine learning, which in turn requires reams of the data that AI learns from, Tesla's advantage in data is likely to translate into a huge advantage in making safe and effective self-driving cars. Indeed, Jonas has argued that Tesla's new mass-market Model 3 sedan could be up to 10 times safer than the average car.

Finally, the Big Five have entrenched themselves in a more traditional fashion, by using their highly valued stock and their enormous amounts of cash to buy other companies, something they've done far more aggressively in recent years. Together, Google, Apple, and Microsoft have roughly a quarter of the cash reserves in the entire S&P 500. Google, the most active buyer, has averaged one acquisition a month. Acquisitions have become increasingly important as a way to gain new technology and new engineering talent, expand into new markets or new product areas, and in some cases squelch potential competition. And since no competitor has the resources to outbid the Big Five, it's another way in which simply being big makes it easier to keep getting bigger.

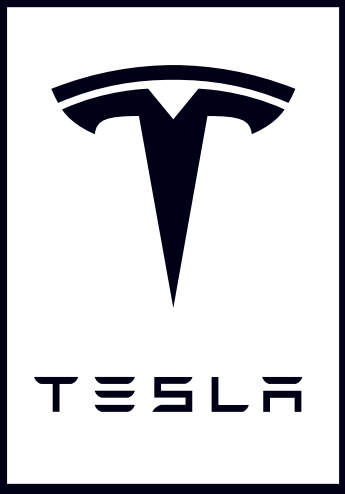


Digital monopolies

So on the one hand, we have a digital economy that, for all the value it's created, has not dramatically improved economic growth or wage growth for ordinary workers; on the other, a big chunk of that economy is dominated by a very small group of players. And what's interesting is that there's some reason to think these two things are, in fact, connected to each other.

To begin with, the most important fact about platform companies is that they scale, meaning that they can create enormous amounts of value while employing a relatively small number of

Tesla & GM



Market capitalization:
$59B

of employees:
17,782

Revenue:
$7B

of cars delivered in 2016:
76,230



Market capitalization:
$52B

of employees:
225,000

Revenue:
$166B

of cars sold in 2016
10,008,000

The Big Five

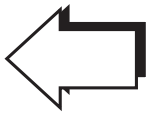
1. *Apple* *$807B*
2. *Alphabet* *$690B*
3. *Microsoft* *$559B*
4. *Amazon* *$484B*
5. *Facebook* *$447B*

MARKET CAPITALIZATION AS OF JUNE 6, 2017

workers. This is a good thing from the perspective of efficiency. But it also helps explain why today's digital giants have a smaller impact on the economy than dominant companies had in the past. Together, the Big Five employ around 400,000 full-time workers in the United States. That might sound like a lot. But roughly half of those workers are Amazon employees, many in relatively low-skill, low-paying warehouse jobs. And it's actually fewer employees than General Motors alone had in 1979, when the U.S. workforce was a lot smaller. What's more, where GM's production led to eight jobs in its supply chain for every one person it employed directly, the ripple effects of the Big Five's businesses, with the exception of Apple, are much smaller. The result is that the rewards of the digital economy are more concentrated among a small number of workers than the rewards of the industrial economy were.

This is aggravated by the fact that the Silicon Valley dream of starting a company in your garage and making it into a huge business has become less realistic than ever. Even though billions continue to be poured into venture funding (more than $200 billion between 2011 and 2016), and even though the number of so-called high-growth startups has not declined in recent years, work by the MIT economists Scott Stern and Jorge Guzman shows that fewer of those startups are succeeding than did in the past. Of course, we still have the Teslas and Ubers (or Lyfts) of the world. But they are rarer than they once were. And one plausible reason is the sheer scale and scope of the Big Five, which can meet competitive challenges either by copying the innovations of others (as Facebook arguably did with Snapchat), and thereby making them seem superfluous, or simply by buying potential competitors at an early stage. Regardless of why it's happening, though, the result is less dynamism in the economy, and less spreading of the wealth.

One obvious solution to the problems caused by the concentration of power in just a few companies is to break up the Big Five or regulate them as, in effect, public utilities. And of late, there have been more and more calls for dramatic action. But this is difficult for a variety of reasons. First, these companies don't, for the most part, fit the stereotype of the monopolist. They aren't "natural monopolies," like power companies, in markets where it would be practically impossible for a competitor to arise. Anyone who wants to build a new search engine, or a new online retailer, can do so. These companies also, with some exceptions, didn't achieve their dominance by engaging in conventionally anticompetitive behavior so much as they exploited the nature of the digital economy to build and maintain their empires.



Nor do you hear too many complaints from consumers, although issues of privacy obviously still matter. Indeed, relative to industries like cable television or airlines, digital companies tend to do well on customer satisfaction, and the digital economy as a whole has become a cornucopia of "free" stuff (paid for with consumers' attention but not their cash). And even though, in practice, consumers often get locked in to the technologies these companies offer (if only because once your data is in the cloud, it's so much easier to stay than leave), the businesses have no real leverage over consumers. Most also continue to pour billions of dollars into research and development, and they are constantly upgrading their products and services. So it's hard to argue that these giants have been anything but beneficial to consumer welfare, which since the 1970s has been the standard that antitrust regulators apply.

One obvious solution to the problems caused by the concentration of power in a few companies is to break up the Big Five or regulate them as, in effect, public utilities. This is difficult for a variety of reasons.



Indeed, when we look at what the digital economy has done over the past two decades, what becomes clear is that it has created an enormous amount of value for consumers and for a small group of big companies, even as it has diminished competition, centralized power, and made life much more difficult for businesses that produce content or try to compete with the economy's dominant players. (In one way or another, if you want to make money in the digital economy, you will almost certainly find yourself working with, rather than against, one of the Big Five.) In the industrial economy, the benefits were spread widely among companies, employees, and consumers. The digital economy is giving us a world in which the benefits are concentrated among consumers and the Big Five who serve them. Everyone else just lives in it.

James Surowiecki is the author of The Wisdom of Crowds *and a senior story producer at* Vice News Tonight.

DATA FROM GM AND TESLA ANNUAL REPORTS

A Job Plan for Robots and Humans

Melonee Wise's startup Fetch Robotics employs about 50 people and more than 125 robots. She intends to create many more jobs for both.

By Tom Simonite



One perk of working for Melonee Wise's startup Fetch is that if your feet are tired you can glide around the office on the back of a squat wheeled robot. More usually, she and her roughly 50 employees keep themselves busy designing, building, and selling the machines to work in warehouses or factories across the globe. The San Jose company's machines are canny enough to work safely alongside people without requiring any changes to a facility—all they need is a map. Wise, one of *MIT Technology Review*'s 2015 Innovators Under 35, built her first robot, Zippy, in college, out of plywood and scavenged computer parts. She spoke with San Francisco bureau chief Tom Simonite.

Do your robots replace people?
No one has ever lost a job because of our robots. Customers need us because they just can't hire enough people. There's 20 percent annual turnover and an estimated 600,000 jobs in the United States going unfilled.

Did I just see a robot go past the door of this meeting room carrying a box of cookies?
That's snackbot. We use our commercial software internally, and you can type into Slack and say "Robot, I'm hungry," and it brings cookies to you. The team uses that because it's fun and you don't have to walk to the kitchen, but it also gets them closer to the product.

Why do you market the versions of your robots with arms and grippers only to researchers, not for warehouse work?
Today it costs about $40,000 to lift up a gallon of milk—that's the cost of a robotic arm that can do that. We're trying to lower the cost—for example, by reducing the number of components in a motor. We're also working on programming the robot by demonstration—basically teaching it a task by hand.

What kinds of relationships do people have with your robots?
Some kind of projection of emotion is inevitable, and you want it to be safe and productive. Making the robot much shorter than a person helps people perceive it more as a child than an adult. A lot of people associate our robots with puppies: they call them pups. We want that, because people are much more inclined to help something that seems simple.

There's a lot of excitement about robots advancing quickly thanks to machine learning. Do you buy it?
I think people are very bad at predicting time lines in the evolution of products. And there are false indicators in the world that give people the perception that things are moving faster than they are. One is facial recognition in things like Facebook. All of that is predicated on a person being centered in the image and well lit. If the lights in this room were very dim, I could probably figure out who you are, but a robot would have no hope in hell. People aren't wrong to get excited. But personally I have a hard time even imagining in my lifetime that we'll have generalized home-purpose robots for cleaning and doing dinner.

When you founded Fetch, you designed and built the first robots. Now that you have customers, has your job changed?
I've become more and more focused on not just making the robots work but making sure that we know where they're going to work. I have been absurdly focused on usability in the last year and a half. For many of the people who interface with our technology on a day-to-day basis in a warehouse, the Web browser is already a challenge. Adding robots into the mix is like dropping an elephant on them.

Not many Silicon Valley founders are women. Is it harder to make it as a female founder?
I haven't lived through a guy's experience, so it's really hard to say. There are only a couple of instances where I've run into some awkwardness around being a woman, most of the time on the subject of what my personal family planning goals are. It's so hard to be a founder anyway—it takes the right person, and that isn't specific to men or women. But some things I perceive as a little bit easier. One is hiring [diverse candidates]. I think women and other minorities don't opt out when applying [to Fetch] because they look at the Web page and see myself and the diversity of the staff that we already have.

What causes the lack of diversity in tech?
How do you blame an industry for being full of men when the pool that they're pulling from is full of men? If you look at my childhood, my dad's side of the family was like, "Yes, do it—get an education." My mom's side was like, "You're wasting your time—you need to go start a family, focus on the things that you are good at because you're a woman." I did pick a side, but that doesn't mean it wasn't hard to hear.

You said no one has lost a job to a Fetch robot, but over time, aren't products like yours going to displace human workers?
There will be displacement. All technology does it. But this is not a technology discussion, this is a socio-political discussion. Technology is going to continue to advance, and technology always creates jobs. Long term, as Fetch deploys more robots, we'll need more people to maintain the robots, to program the robots, to install the robots. Those jobs are coming. The question is, will the people whose jobs are displaced be retrained to do the new jobs that are available? Typically, training is motivated through social, political, and economic means—incentivizing companies or making college free. I personally believe that we as a country should move toward universal basic income.

How would that be different from existing state support?
I was on welfare most of my childhood, and it didn't get you everything. I still had to stand in food lines and go to food pantries. I think of universal basic income as saying we all have the right to live our lives and experience the basic standards of life—and anything above that is yours to get.

ANDREW PAYNTER

ZIP







MIT TECHNOLOGY REVIEW
VOL. 120 | NO. 4



LINE

Help from Above

In Rwanda, an early commercial test of unmanned aerial vehicles is delivering urgently needed medical supplies.

By Jonathan W. Rosen

Photographs by Jason Florio

Previous Pages

A Zipline technician uses a catapult mechanism to launch a drone carrying blood sachets toward Kabgayi Hospital. Once a drone is over the hospital, its package drops to the ground via parachute.

YOU

can hear the drone before it's visible, whining like a mosquito above the hillside grounds of Rwanda's Kabgayi District Hospital. Emerging through a patch of fog, roughly 100 feet in the air, the small plane quickly disappears again, circling in an oblong pattern as it descends toward an altitude low enough to make its drop. After a period of silence, it's suddenly back, swooping over the roof of Kabgayi's accident ward to drop its payload on the driveway with a thud. On the ground lies a red cardboard container, roughly the size of a shoebox, attached to a parachute made of wax paper and biodegradable tape. The contraption may resemble a children's art project, but its contents are lifesaving. Packed tightly inside are two units of human blood, which will probably soon be used for transfusions during surgeries or complicated childbirths, or to treat young victims of malaria.

The plastic sachets of blood are among the first commercial products ever delivered by drone, part of a partnership between the Rwandan government and the Silicon Valley–based robotics firm Zipline, which began introducing the blood drops at Kabgayi in late 2016. The service, which is now delivering to seven of 21 planned facilities, is still in its infancy. Yet it has already had an impact. In the past, hospital staff would make three drives per week to procure blood products in the capital, Kigali, 60 kilometers away, a three- to four-hour round trip. Emergencies meant additional trips—sometimes resulting in life-threatening delays. Now, a Kabgayi lab technician simply taps out an order on a smartphone and Zipline's distribution center, located five kilometers from the hospital, will have a drone there within 15 minutes. "Before, it was a serious problem to have blood when we needed it," says Espoir Kajibwami, a surgeon who served as Kabgayi's medical director until February. In emergency cases, he says, the hospital would often send the patient to the national referral hospital in Kigali rather than wait for the blood to arrive.

Zipline's blood deliveries come at a time of great activity in the world of drone-enabled commerce. Last December, three years after announcing its much-publicized Prime Air service, Amazon conducted its first commercial drone delivery, to a farmhouse in rural England. The month before, the convenience-store chain 7-Eleven completed 77 on-demand drone deliveries of pizza, Slurpees, and over-the-counter medicines to customers in Reno, Nevada. UPS, which has helped finance Zipline's operations through a $1.1 million grant from its charitable foundation, delivered a package in February with a drone launched from the top of one of its signature brown trucks. Flirtey, the drone maker behind the 7-Eleven pilot, has also tested the delivery of medicines in understocked parts of Appalachia. Another U.S. firm, Matternet, has conducted test flights in collaboration with UNICEF to deliver infant HIV test kits in Malawi.

Yet Zipline, which uses fixed-wing drones that have a greater range and are more resilient in bad weather than the more common multicopter models, is the first in the world to offer regular delivery of emergency medical products.

Founded in 2011 under the name Romotive, the company first gained notoriety as the maker of Romo, an iPhone-powered robotic pet, before CEO Keller Rinaudo decided to seek a product that would have a greater social impact. Soon he and Zipline cofounders William Hetzler and Keenan Wyrobek were scouring the developing world to learn how drone-based logistics could help save lives.



Lab technician Prosper Uzabakiriho collects a package for Kabgayi Hospital containing blood dropped by paper parachute from the Zipline drone.

On two separate trips to Tanzania in 2014, Rinaudo and Hetzler, who'd first worked together as Harvard freshmen to build a rock-climbing gym on campus, met a researcher working with a local NGO who'd created a text-message-based health surveillance system. Through a network of community health workers, the project tracked hundreds of health emergencies across the country, including snakebites, possible rabies cases, and severe postpartum hemorrhaging. Frequently, however, reaching remote patients with treatment in time was cost-prohibitive or logistically impossible. "You have this database full of tragic human stories," says Hetzler, 30. "The obvious missing piece was a way to very rapidly respond to that demand and get the product into a place where it isn't otherwise readily accessible." The pair came away from their visits feeling they could build that.

Under the technical direction of Wyrobek, Zipline's engineers got to work at the company's headquarters in Half Moon Bay, California, developing a craft with twin electric motors that could carry a 1.5-kilogram payload and operate in almost any weather. As the technology progressed, Rwanda emerged as an ideal spot to test the Zipline vision. Similar to Tanzania, its neighbor to the east, the small East African country is home to a predominantly rural population that is

A Zipline technician carries one of the company's drones.

The drone is monitored as it follows its pre-programmed flight path.

plagued with transportation challenges. A majority of Rwanda's 478 health centers, and many of its 35 district hospitals, can be accessed only via poorly maintained unpaved roads, which often spiral into its famed "thousand hills" and are difficult to traverse by vehicle, particularly during the twice-yearly rainy season. Unlike Tanzania, though, Rwanda is compact: with 12 million people in an area the size of Maryland, it is the most densely populated country in mainland Africa. This meant that Zipline's drones, which have a flight range of 150 kilometers, could serve nearly half the country from a single launch site.

Zipline's idea caught on with Rwandan authorities, including the country's civil aviation body, which altered regulations to enable its drones to operate. In mid-2016, Zipline signed a deal with the Rwandan government to build a distribution center near the town of Muhanga.

This hilltop site, erected on a former maize field, is now known as the nest: a fenced-in plot consisting of a white circus-style tent housing a blood storage facility, 13 drones (nicknamed "zips"), and a small staff of young Americans and Rwandans. On one side of the tent, two stainless-steel launchers, facing opposite directions to account for changing winds, employ a system of bungee pulleys to catapult the 12-kilogram drones into the air at 84 kilometers per hour. On the other, two brown inflatable mats cushion the zips' landing on return. Once the drones are airborne, cruising over an undulating landscape dotted with banana trees, cassava fields, and tin-roofed houses, an operator monitors their path on an iPad, staying in constant touch with air traffic control in Kigali. All routes, developed using a 3-D satellite map followed by detailed manual ground surveys, are pre-programmed using real-time kinematic satellite navigation, which—along with an inertial navigation system—enables the payload to drop within a target area five meters in diameter. "Accuracy is extremely important," says Hetzler, adding that if the drops were less precise, packages could end up on roofs, in trees, or in other inaccessible spots that could destroy the operation. He says the company is developing technology that will automate the ground surveying process.

Zipline's plans for Rwanda include scaling up to a wide range of medical products, including emergency rabies vaccines; drugs to treat HIV, tuberculosis, and malaria; contraceptives; and diagnostic test kits.

Yet blood represented a natural starting point. After all, it has a shelf life of only 42 days, must be kept refrigerated, and is frequently needed on an urgent basis.

Today, Rwanda's Ministry of Health stores blood at a national center in Kigali and four regional depots around the country; its 58 facilities equipped to handle blood transfusions, mainly hospitals, keep a small inventory of common blood types and must continually restock from the depots or national center. Sometimes the stock on hand is sufficient for emergency transfusions. Often it is not, particularly if the patient has a less common blood type. In these cases, the facility must refer the patient or collect the blood by car, motorbike, or truck.

Among the individuals most vulnerable to blood delivery delays are pregnant women. Although Rwanda's maternal mortality rate has declined by more than two-thirds since 2000, an improvement linked to increases in contraceptive use and a move away from home births, complications of childbirth remain a leading cause of death. The World Health Organization estimates that Rwanda has one maternal death for every 344 live births, 20 times the rate in the United States and 97 times the rate in the top-performing countries in Europe. More than half of maternal deaths occur after childbirth, and 26 percent are the result of hemorrhaging. Faster, more reliable access to blood could help reduce this number, along with lives lost to accidents and malaria-induced anemia, which is common in small children.

Kabgayi, one of Rwanda's largest district hospitals, faces a higher blood demand than most. The colonial-era facility, several aging brick buildings on the site of a former Catholic mission, uses up to 100 units per month, according to Kajibwami, the for-





Transporting water near Kigali.



mer medical director. Roughly half of that demand comes from the maternity ward, which recorded more than 4,600 births in 2016. The hospital's high volume of patients was evident on a February morning from the crowd that sat on wooden benches in its open-air reception area: women in colorful patterned dresses, men in fraying jackets, uniformed schoolchildren, and even a small group of prisoners dressed in pink and orange smocks.

Although many Kabgayi patients come from hard-to-reach surrounding villages, the hospital itself, which sits just off a well-maintained paved road to Kigali, isn't difficult to access. The facility is also too close to Zipline's distribution center for drone deliveries to make much sense absent plans to scale up to other hospitals: if its drones failed to function properly, the company reckoned when selecting Kabgayi as its inaugural site, it could still deliver blood quickly by truck. Yet most of the facilities Zipline plans to service are considerably more remote. Personnel at Muhororo District Hospital—Zipline's second pilot facility, which received its first drone delivery at the end of February—could spend an entire day battling rutted dirt roads, and the occasional landslide, to make a blood collection in Kigali.

Zipline's impact is hard to quantify. Difficulty accessing blood is not the only factor inhibiting effective care: physicians and other skilled personnel are in



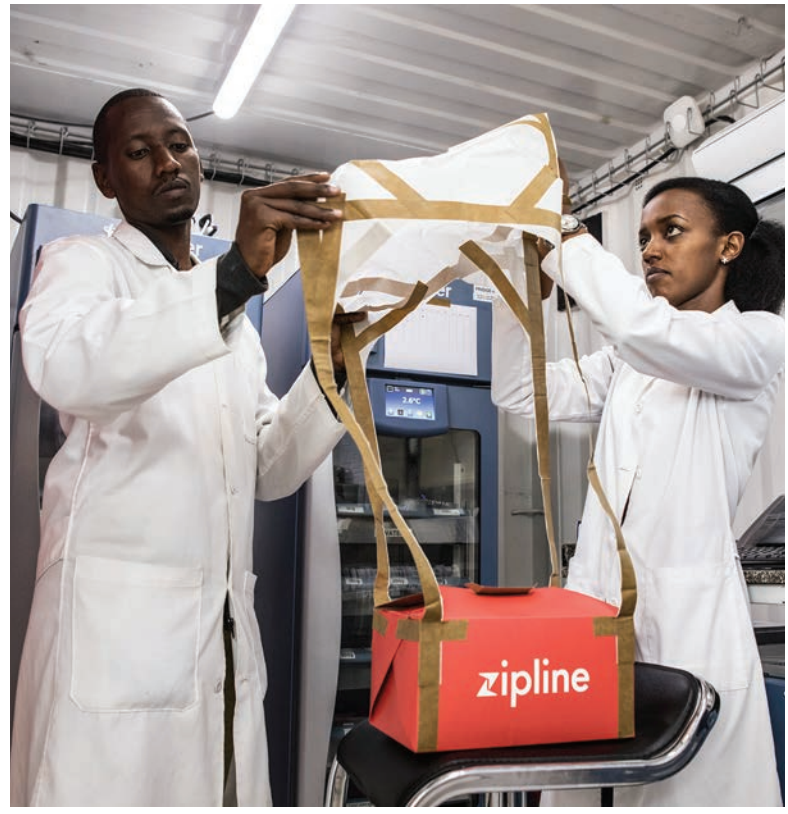
ACCURACY IS EXTREMELY IMPORTANT

A Zipline technician checks a drone during preparations for flight.






CNTS RWANDA - KIGALI
O +


zipline


CONSULTATION
AHO
BASUZUMIRA
AHO MUGANGA ASUZUMIRA N°2

L: Prosper Uzabakiriho inspects a blood sachet after retrieving it in its cardboard delivery box from the driveway of Kabgayi Hospital.
R: After receiving a blood request, Zipline staff check the blood, its packaging, and a parachute made of wax paper and biodegradable tape before the package is loaded onto a drone for the five-kilometer flight to the hospital.

The open-air waiting room at Kabgayi, one of Rwanda's largest district hospitals.

short supply, and many patients have trouble reaching medical facilities. A majority of pregnant women plan to deliver at less equipped, more remote health facilities that typically lack doctors or blood transfusion capabilities. Complicated births, therefore, are usually referred to hospitals—a process that can cause critical delays in treatment for which drones have no answer.

Few inside Rwanda, a highly authoritarian state, dare to openly criticize the government or its programs. Yet many who've heard about the project are skeptical in private and question why authorities are investing in unproven high-tech gizmos when demand for fundamentals like ambulances and health workers still exceeds supply. To some critics, the project's most important benefit could be the good PR it generates, which distracts from the Rwandan government's spotty human rights record and the acute levels of poverty in much of the country. Rwanda, which is highly reliant on foreign aid, has long marketed itself as a regional hub of tech and innovation—an idea that remains more of an aspiration, and a strategy for generating outside interest, than a reflection of current reality.

It doesn't help that the project's costs are not disclosed. Rwanda's government pays Zipline a fixed price per delivery, with a minimum volume guarantee, but neither will divulge the terms. Hetzler admits that deliveries by drone, for now, are more expensive than standard trips by road. Still, he argues, this cost will be offset if Zipline makes possible a far more agile supply chain—for blood as well as more durable products like medicines—in which fewer items must be kept at last-mile facilities, minimizing waste and more effectively preventing stock-outs. As a general rule, per-delivery costs should decrease as flight volume expands. Bruce Y. Lee, a professor at the Johns Hopkins Bloomberg School of Public Health who led a research team last year that simulated a drone-based vaccine supply chain in Mozambique, says drones "certainly have the potential to both decrease costs and increase demand fulfillment," depending on the frequency of flights, the distances traveled, and the speed of trips on local roads.

In Zipline's case, better logistics also mean that facilities can access products with shorter shelf lives or unique storage requirements. In addition to blood units of all types, Zipline's existing service offers platelets, fresh frozen plasma, and cryoprecipitate, all of which promote blood clotting and were previously underutilized because they were too difficult to store at health facilities.

A fuller picture of Zipline's impact on Rwanda will probably not take shape for some time. After commencing operations at Muhororo and five additional sites, the company is working to begin deliveries at the 14 other planned facilities by the end of this year. In addition to introducing new products, Zipline plans to build a second distribution center in the country's east. This, together with future-generation drones expected to have longer ranges, would enable it to cover all of Rwanda except Kigali, where products are stocked in sufficient quantities already.

Additional countries are also in the pipeline. According to Hetzler, a "long list" of governments have expressed interest in Zipline's technology, and the company has already etched agreements with Tanzania's Ministry of Defense and Civil Aviation Authority that would enable it to launch there next. To finance its expansion, Zipline raised $25 million last summer, bringing its total capital raised to $43 million. Still, it is choosing its next locations carefully, with a priority given to countries that offer a chance for social impact and are willing to make necessary regulatory changes and a firm financial commitment. Eventually, Zipline also aims to operate in wealthier countries, including the United States, where regulations currently restrict drone flight to within an operator's visual line of sight absent special review.

Back at Kabgayi Hospital, health workers say the service has already helped. Inside the modern laboratory building, where blood is refrigerated in a small, ceramic-tiled room, Prosper Uzabakiriho, a white-coated lab technician, says that cutting back on lengthy road trips to Kigali has allowed many patients to be treated more promptly and freed up time for staff to perform their duties. At the time of my visit in mid-February, the

Boys play soccer on a hilltop pitch. Kigali is in the background.



facility had received an average of 11 deliveries per week since the start of the year and was entirely reliant on Zipline for blood products during normal business hours. By late May that had climbed to 20 per week, including weekend deliveries. Rwanda's civil aviation authority has not yet given it the go-ahead to operate at night, so Kabgayi will continue to rely on its former blood supplier in Kigali for emergencies after hours. (According to the Ministry of Health, the country's five existing blood depots, which also serve as centers for blood donation, will remain in operation even as Zipline scales up.)

The surgeon Kajibwami credits the service with helping him respond to one particularly urgent case: a woman who began hemorrhaging after surgery to remove an ectopic pregnancy. The ensuing operation, which entailed a partial hysterectomy, required six units of blood, more than Kabgayi had in stock, and thus an urgent call was made to Zipline. The doctor cannot say whether the patient would have died in other circumstances. Still, he's grateful for the timely delivery. "It was very bad," he says of her condition. "It would have been very difficult to manage it."

Jonathan W. Rosen is a journalist reporting from sub-Saharan Africa. He wrote about e-commerce in Ghana for the January/February issue.



Q+A

Ya-Qin Zhang

The President of Search Giant Baidu Has Global Plans



Ya-Qin Zhang expects investments in artificial intelligence and other technologies to help turn his company into a global leader.

By Tom Simonite

Google and Facebook are household names around the world. Baidu? Not yet. Ya-Qin Zhang, president of China's leading search business, says Chinese companies can become worldwide Internet powerhouses too. One of the biggest fish in China's market of 730 million Internet users, Baidu is trying to open new revenue streams both domestically and abroad by investing heavily in artificial intelligence.

The company employs more than 1,700 AI researchers, including some at a Silicon Valley research center opened in 2014, and was chosen by the Chinese government to run a new national lab intended to make the nation more competitive in machine learning. Zhang even predicts that the self-driving cars Baidu is developing might be in widespread public use before those introduced by U.S. competitors.

Domestically, revenue growth in its existing search and ads business wilted in 2016 as smartphone sales slowed in China and rivals including Tencent, owner of the popular mobile messaging and e-commerce platform WeChat, closed in.

Baidu's reputation and ad business were damaged as well by a scandal in which a student died from cancer after a sponsored result displayed by the company's search engine led him to expensive and questionable treatment.

Zhang, who previously led Microsoft's research and business operations in China, took time out from meeting staff at Baidu's U.S. R&D center in Silicon Valley to meet with *MIT Technology Review*'s San Francisco bureau chief, Tom Simonite. An edited version of their interview is below.

You say that Baidu already has more than a dozen mobile apps used by 300 million people outside China every month. What is your strategy for other markets?

We are very focused on mobile. Our apps are popular in India, Brazil, parts of Southeast Asia like Indonesia. Those are places where the market is very similar to China, but maybe [China as it was] a few years ago—large populations that are mobile-centric. And that gives an advantage over U.S. companies.

Will you offer products in the U.S.?

While we have over 20 million users in the U.S., our focus is on R&D and talent. Right now PhDs from the best Chinese universities are on par with the top U.S. graduates. We have a lab here in the U.S. because we also need more senior talent. The AI team in the U.S. has already developed a new speech-recognition system that is playing a role in our product.

You have spent a lot on artificial-intelligence research. How does that translate into new revenue?

AI will be like electricity, in that it will drive a lot of new business to many industries. We have built a platform we and partners can build on to offer services in anything—health care, education, logistics. As one example, we have a new partnership with Bosch around using our high-definition maps for autonomous and semi-autonomous cars. Many of these [services] will probably be bigger than search, which will continue to be our core business for many years and is also being transformed by AI.

Your chief scientist Andrew Ng, who previously helped establish Google's machine-intelligence research, left the company recently. Is this a long-term setback?

Andrew is a great guy, and it's a loss to Baidu here in the U.S. and to the wider company. But he built a large AI team with a lot of talent, and Baidu has a very deep bench of leaders. Even in the U.S. labs, more than 60 percent of employees—on autonomous driving, security, our advertising platform—worked for somebody else.

Why would top talent in artificial intelligence want to work for Baidu's U.S. lab over that of an American company?

We allow researchers and engineers to make a bigger impact with their work. You get to be part of the two most exciting and dynamic technology ecosystems on Earth: Silicon Valley and the Internet economy of China. The Chinese market gives researchers access to a massive amount of data and a consumer base incredibly quick to adopt new technology.

Baidu is developing and testing autonomous cars in both the U.S. and China. The leading projects in this area are almost all based in the U.S. Where will this technology be commercialized first?

In China the roads and traffic are more complex than in the U.S. But the first batch of applications will be in restricted areas, and it could be easier in China because regulators are more flexible and open-minded. Mayors of many different places want us to try on their roads and are willing to build out special areas for us. A couple of major ports in China want to use our technology on trucks.

Are there areas where you are ahead of U.S. competitors?

I think we deploy new algorithms into our products much faster. The Chinese market is hypercompetitive, and to avoid being beaten by competitors, the cycle of [translating new technology into products] is much shorter than in places like the U.S. Startups in China are even faster than us at making new AI algorithms into products, and that's something to look out for.

How are Chinese tech companies doing overseas so far?

We're all just beginning. None of the biggest Chinese Internet companies has proven success outside of China, but I think in the next five years, 10 years, you will see those companies become global. Alibaba, Tencent, Baidu, and others—we are all deploying teams and building up talent. Next will come more products and many more users.

BAIDU

50 Smartest Companies 2017

It Pays to Be Smart **By David Rotman**

Superstar companies are dominating the economy by exploiting a growing gap in digital competencies.

Our economy is increas- ingly ruled by a few dominant firms. We see them everywhere, from established giants Amazon, Facebook, Google, Apple, and Walmart to fast-growing newcomers like Airbnb, Tesla, and Uber. There have always been large companies and outright monopolies, but there's something distinctive about this new generation of what some economists call superstar companies. They appear across a broad range of business sectors and have gained their power at least in part by adeptly anticipating and using digital technologies that foster conditions where a few winners essentially take all.

Our annual list of the 50 Smartest Companies includes many of these firms, but it's not merely a list of today's biggest or most profitable players. It highlights technologically innovative companies whose business models allow them to exploit these advances. The list is our best guess as to which firms will be the dominant companies of the future. Amazon and Facebook and Google are on it, but so are plenty of newcomers. Though they might be unfamiliar to you today, we believe they have an inside track to take advantage of the technologies, such as artificial intelligence, that will define business in the coming years.

Being smart about innovation won't guarantee that these firms become superstars. But it does, at least, give them the potential to create and dominate new markets in an increasingly competitive business environment.

The emergence of superstar companies has, in many ways, helped to define our era. Digital giants, in particular, have cleverly leveraged the Internet, so-called network effects, and big data to become hugely profitable while providing indispensable services—like free Web search and easy online shopping—and devices that have changed our lives (see "Why Tesla Is Worth More Than GM," page 28).

But Internet companies aren't the only ones to become superstars. According to recent research by economists at Harvard and MIT, the share of sales by superstar companies—which the authors define as the four largest firms in a given industry—has gone up sharply in all the sectors they looked at, from transportation to services to finance. The trend toward superstar firms is accelerating, says Lawrence Katz, a Harvard economist and coauthor of the study. It has become more uniform across industries and developed economies during the past decade or so. These companies' dominance is particularly strong in markets undergoing rapid technological change. Katz says that's probably because of the wide disparity in how well companies take advantage of new advances. In other words, you have to be the smartest company in your field or you might as well not bother.

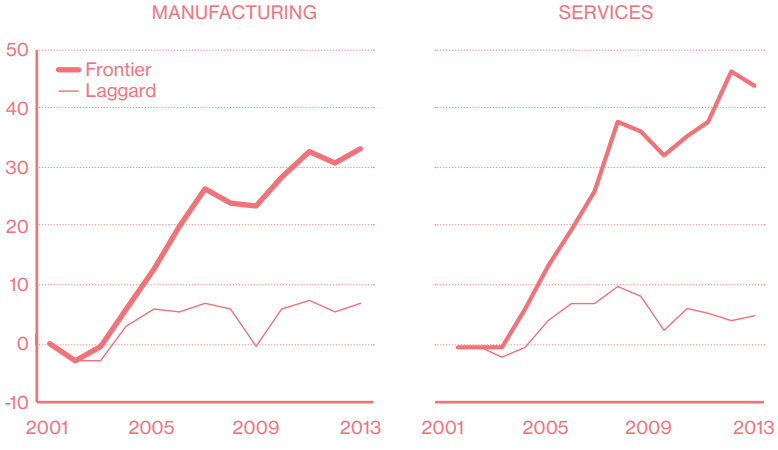
In itself, that might not be bad. But the authors identified a deeply troubling result of an economy where just a few top-tier companies dominate. One of the economic truths of much of the 20th century was that the portion of the country's overall income that went to labor was constant; as the economy grew, workers got a proportionate share of that growing pie. But labor's share of the national income has been shrinking over the past few decades. This is true in many countries, and the decline speeded up in the United States in the 2000s.

The trend puzzles economists. Some suggest it reflects the rise of cheap robots that can do the jobs of human workers, but the data isn't con-

vincing. Instead, Katz and his coauthors blame the emergence of the superstar companies. As these companies grow and become more efficient and more adept at using digital technologies, they need fewer workers relative to their soaring revenues. The fact that these labor-frugal firms have so much of the market share in their sectors means labor gets a smaller portion of the nation's overall income.

Compounding the problem is that superstar companies, which desire the best possible talent, tend to pay much better than anyone else. This dynamic is deepening the divide between the country's economic winners and losers. Nicholas Bloom, an economist at Stanford, and his colleagues have shown that about one-third of the growth in U.S. income inequality since 1980 can be explained by the disparity between the pay premiums of a few elite companies and the salaries most workers earn. Fewer and fewer people—mostly a select group of highly trained professionals—are enjoying the vast profits generated by these top companies.

THE WIDENING GAP BETWEEN FRONTIER FIRMS AND LAGGARDS CAN BE SEEN IN THE PERCENTAGE INCREASES IN LABOR PRODUCTIVITY SINCE 2001



Frontier firms are the 5 percent of companies with the highest productivity within each industry; laggards are the rest of the industry. Data is from 24 OECD countries.

D. ANDREWS, C.CRISCUOLO, AND P.N. GAL (2016) "THE BEST VERSUS THE REST"

It is "certainly a big part of the [economic] anxiety" that is plaguing the country, Katz believes.

The rise of the superstar companies also might help explain another disturbing economic trend. Despite the proliferation of impressive new advances in software, digital devices, and artificial intelligence over the last decade and the great profits generated by Silicon Valley, economic growth in the United States and other developed countries has been sluggish (see "Dear Silicon Valley: Forget Flying Cars, Give Us Economic Growth," July/August 2016). In particular, an economic measure called total factor productivity, which is meant to reflect innovation, has been dismal (see "Tech Slowdown Threatens the American Dream," May/June 2016). How can overall growth be so lackluster while the high-tech sector is booming?

Economists with the Organization for Economic Cooperation and Development think they have found the answer. It turns out that productivity at the top companies in various sectors—what the OECD economists call the frontier firms—is growing robustly. These are the companies making the best use of the Internet, software, and other technologies to streamline their operations and create new market opportunities. But most companies aren't actually harnessing new technologies very effectively. And the relatively poor productivity of these laggards, says OECD economist Dan Andrews, is dragging down the overall economy. "Technologies are increasingly complex, and many firms may lack the competencies to adapt," suggests Andrews, coauthor of the OECD study, which looked at the United States and 23 other developed countries.

In some ways the OECD findings are encouraging, because they demonstrate that recent innovations do—in the hands of top companies—have the potential to strongly improve productivity. But surprisingly, says Andrews, the laggards seem to be making little progress toward catching up; new ideas and business practices aren't trickling down as rapidly as they should. The reason isn't entirely clear, he says. But it seems that the economy is less dynamic and efficient at "dispersing" new technologies than we might think.

Such findings help drive home the importance of the 50 Smartest Companies list. Be assured, there are no laggards on it. But the research by Andrews and others also shows why we need a better business climate—one that allows more startups and fresh ideas to thrive. Today's giant companies are pulling ahead, and a dwindling number of individuals are reaping the financial rewards. There is nothing inevitable about that trend. The advent of complex technologies such as artificial intelligence, which will be critical to future business success and are tricky to understand and master, could widen the gap further. They could also provide ample opportunities for new companies to create markets that don't even exist today. We do need companies to aggressively push the frontiers of innovation. Still, as we celebrate our 50 Smartest Companies, it is worth keeping in mind the importance of distributing know-how, and the wealth it produces, more broadly.

50 Smartest Companies **The List**

1. Nvidia

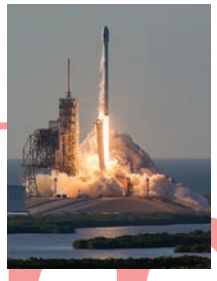
Still improving its chips, originally developed for gaming, to help develop breakthrough technologies like deep learning and autonomous driving.

$3 billion: *spending on R&D to create its new chip*

2. SpaceX

Changing the economics of space travel with its successful landing of rockets to be recycled for multiple trips.

10 percent: *discount being considered for customers who agree to use reused rockets*



3. Amazon

Creating an AI-powered store of the future with Amazon Go while expanding assistant Alexa into phones, cars, and more.

12,000: *number of computer programs that developers have published for Alexa*

4. 23andMe

Vindicated this year when the U.S. FDA granted permission to tell customers whether their DNA puts them at higher risk for some diseases.

1 million plus: *number of customers who have consented to have their genetic information used for scientific research*

5. Alphabet p58

Continues to dominate research into AI while expanding innovation in smartphone software, virtual reality, and self-driving cars.

40 percent: *amount of energy the company says it saves applying machine-learning algorithms from its DeepMind subsidiary to cooling its data center*

6. iFlytek

Its voice assistant technology is the Siri of China, and its real-time portable translator puts AI to remarkable use, overcoming dialect, slang, and background noise to translate between Chinese and a dozen other languages with surprising accuracy.

70 percent: *iFlytek's share of China's market in voice-based technologies*

7. Kite Pharma

Nearing FDA approval of its experimental immunotherapy that uses a patient's own blood cells to combat cancer.

31 percent: *proportion of study participants who were very sick with lymphoma yet showed no sign of the disease six months after a single treatment with Kite's therapy*

8. Tencent

Turning its insanely popular chat platform WeChat into a virtual operating system featuring mini programs.

50 percent: *proportion of WeChat's 770 million daily users on the service at least 90 minutes a day*

9. Regeneron

Biotech with a strong drug pipeline and track record is treating eye and other diseases, and testing treatments for rheumatoid arthritis, asthma, and pain.

500,000: *number of U.K. volunteers whose genetic data it is helping sequence*

10. Spark Therapeutics

Its blindness treatment could be the first gene therapy approved in the U.S. to treat an inherited disease.

1 in 30,000: *estimated number of individuals affected by the disease, Leber hereditary optic neuropathy*

11. Face++

Pioneering new uses of face recognition technology, from helping to solve fraud to "smile to pay."

106: *maximum points on a person's face that its technology tracks*

12. First Solar

Making advances in cadmium telluride cells; building three of the five largest solar projects in the U.S.

$2.9 billion: *estimated 2017 revenue*

13. Intel

Acquisitions in computer vision and AI show it's serious about adapting to new technology.

46 percent: *portion of revenues derived from areas beyond PC chips*

14. Quanergy Systems

Its solid-state version of lidar is cheaper and more compact than conventional versions of a technology essential to autonomous driving.

$250: *price of its S3 lidar sensors for autonomous vehicles*

15. Vestas Wind Systems



Overtook General Electric to become the biggest U.S. installer of new wind power last year and is investing in energy storage.

14: *consecutive number of profitable quarters*

16. Apple

Minting money selling its popular mobile phones and laptops while adding impressive names to its AI research team and promising to do more manufacturing in the U.S..

$257 billion: *cash on its balance sheet, more than the entire market value of General Electric*

17. Merck

In a first, the FDA has approved its immunotherapy Keytruda to treat cancers on the basis of the tumor's genetic characteristics, not its location in the body.

$39 billion: *estimated 2017 revenue, buoyed by sales of Keytruda*

18. Carbon p62

Its novel 3-D-printing process permits fabrication of parts out of a wide variety of plastics.

100,000: *number of pairs of shoes Adidas will print by the end of 2018 using Carbon technology*

19. Desktop Metal

With nearly $100 million from VC firms, GE, Alphabet, and others, this startup is focused on cheap, fast 3-D printing of metal parts.

$120,000: *cost of its first product, to begin shipping in September*

20. Ionis Pharmaceuticals

RNA drug approved for a rare disease, spinal muscular atrophy.

36 plus: *number of RNA-targeted drugs in development*

21. Gamalon

Its technology can write and rewrite its own code, algorithms that will accelerate machine learning.

100 times: *its technology's efficiency advantage over other machine-learning methods*

22. Illumina

After a drop in sales last fall, unveiled a new machine, NovaSeq, that will be capable of sequencing 48 entire human genomes in two and a half days — and could one day push the cost of genome sequencing down to $100.

$850,000: *price of the cheaper of its two NovaSeq models*

23. Facebook

Despite controversies over fake news, live streaming video, and discriminatory advertising, and poor sales of its Oculus VR headset, it continues to work on interesting applications of AI and VR, and its Instagram business is singing.

20: *number of natural-language data sets built into its AI research tool ParlAI*

24. Udacity

Has found a business model for MOOCs by working with corporations to make course material relevant to jobs; now connecting companies to students and graduates, too.

15: *number of "nanodegrees" the company offers in skills for selected jobs*

25. DJI



Ongoing innovation in consumer drones and now expanding into drones for enterprise as well.

50 percent: *estimated North American market share*

26. MercadoLibre

Runs the largest online market in Latin America and is expanding to mobile point-of-sale transactions. Its MercadoPago online payment tool lets users deposit cash into their accounts.

182 million: *number of registered users, a 20 percent increase over the previous year*

COURTESY OF SPACEX; COURTESY OF INTEL; COURTESY OF VESTAS WIND SYSTEMS; COURTESY OF DJI

27. Microsoft

Its fast-growing cloud business has reduced the software giant's reliance on PC sales. Its expanding team of quantum computing experts is working on commercially viable products to compete with efforts by Google and IBM.

$15 billion: *projected annual revenue for its commercial cloud business*

28. Rigetti Computing

Though a startup, it's got its own fab in the Bay Area and an ambitious approach to quantum computing that combines hardware and software, focusing on design that can be easily commercialized.

$64 million: *venture funding raised by the company in the past year*

29. Kindred AI

Combining strengths of humans and robots into exoskeleton suits in a bid to help people and machines work together.

Immersive teleoperation: *the type of technology the company makes, in which a human controls a robot via a wearable device*

30. Sophia Genetics

Evangelists of data-driven medicine are sorting through DNA sequences with AI algorithms to accelerate diagnosis in oncology, cardiology, and more.

106,000: *number of patients tested to date*

31. Tesla

Autopilot limitations, car maintenance problems, and concerns about its solar strategy and ability to produce enough cars have hurt, but cofounder Elon Musk continues to take big bets. Battery cell production has begun at his giant Nevada "gigafactory."

400,000 plus: *number of preorders for its lower-cost Model 3*

32. Oxford Nanopore p 66

Twelve years and $200 million in the making, its inexpensive, portable genetic analyzer has been successfully tested from Antarctica to space and shows promise for on-the-spot diagnostic testing, germ monitoring, and more.

882,000 letters: *record length of a single DNA strand read continuously by one of its machines*

33. Foxconn

Acknowledging the direction of Chinese manufacturing by shifting from low-cost human labor to add extensive robotics to its factories.

60,000: *number of jobs automation eliminated at a single Chinese factory*

34. M-KOPA

Its pay-as-you-go solar power model works well in its African target market, and the company is expanding sales with local communications leader Safaricom.

500,000: *number of homes connected as of this spring*

35. ForAllSecure

Still in startup mode but has garnered attention since last August, when it won the Pentagon's DARPA contest with a bot designed to autonomously spot, test, and fix software security flaws.

14: *number of vulnerabilities one experiment discovered in networking devices*

36. Flipkart

Benefiting from the consolidation of India's competitive e-commerce sector, including $500 million investment from eBay.

$11.6 billion: *company's current valuation, the highest for any Indian e-commerce startup*

37. Bluebird Bio

Leading gene-therapy company focused on engineered T cells that recognize and kill cancer and other conditions. Its treatment for sickle-cell disease appears promising.

66: *percentage increase in stock price over the past year*

38. Adidas



Commercial production at its robot-heavy factory in Ansbach, Germany, to begin this year, producing locally and on demand. A second factory has been announced in Atlanta.

300 million: *number of pairs of shoes Adidas makes each year, largely in Asia*

39. IBM p 72

Exploring new technologies like blockchain and cloud AI while continuing work on important long-term challenges like quantum computing.

400: *number of customers the company has worked with on blockchain applications*

40. General Electric p 78

Moving to incorporate AI into its businesses as it focuses on technological innovation in wind and renewable energy, data-driven services, and other business lines.

60,000: *number of jet engines GE says will be connected to the Internet by 2020*

41. Alibaba

Quickly expanding the artificial intelligence in its Alibaba Cloud platform, including industry-specific products, and launching a global electronic trade platform to build its business with small and medium-size companies around the world.

57 percent: *Alibaba's share of Chinese online commerce*

42. HTC

Despite some executive turnover and a tough year financially, the company has interesting prospects in virtual reality and access to China's enthusiastic VR users.

1,500: *number of pieces of content outside developers have created for its VR system, Vive*

43. Blueprism

Its software helps companies including banks and insurers use AI to do back-office clerical tasks.

189: *number of deals signed by the company in 2016, more than four times its 2015 number*

44. Africa Internet Group

The online platform has consolidated all its consumer Web services—shopping, travel, food delivery, real estate sales, car rentals—under one name, Jumia.

500,000: *number of African companies that use Jumia's platform*

45. Veritas Genetics

Launched by well-regarded researchers and funded with venture capital backed by Chinese and U.S. pharmaceutical companies, it will sequence anyone's genome for just $1,000 and interpret it too. In 2017, started offering to sequence newborns in China.

1,250: *number of conditions, risks, and traits it will tell parents about in their newborns*

46. Daimler

Delivering the first run of its short-haul all-electric truck this year while working on vehicle connectivity and autonomous driving for cars.



200 kilometers: *range of its lithium-battery-powered eTruck*

47. Salesforce

Looking to push AI-generated tools, such as an algorithm designed to summarize documents, to its massive user base.

20 percent: *market share in customer relationship software*

48. Snap



Though its glasses have not been a runaway hit, it is succeeding with entertainment through its filters and pioneering the use of machine vision and augmented reality for socializing.

3 billion: *number of snaps users create each day*

49. Ant Financial

Its Alipay business, dominant in China, is eyeing new markets, including India and Korea. Meanwhile, it's exploring AI as a means of underwriting loans, and blockchain as a method of improving security in back-office operations.

450 million: *number of annual active users*

50. Baidu

Despite the high-profile loss this year of its well-regarded head of AI, Andrew Ng, the company is doing important work in the field, including leading China's new national deep-learning lab.

1,700: *number of employees dedicated to working on AI*

COURTESY OF MICROSOFT; COURTESY OF ADIDAS; COURTESY OF DAIMLER; COURTESY OF SNAP

5. Alphabet

The company thinks it can wrest the cloud computing market away from Amazon by helping companies make use of machine learning.

By Tom Simonite

Early in 2015, artificial-intelligence researchers at Google created an obscure piece of software called TensorFlow. Two years later the tool, which is used in building machine-learning software, underpins many future ambitions of Google and its parent company, Alphabet.

TensorFlow makes it much easier for the company's engineers to translate new approaches to artificial intelligence into practical code, improving services such as search and the accuracy of speech recognition. But just months after TensorFlow was released to Google's army of coders, the company also began offering it to the world for free.

That decision could be seen as altruistic or possibly plain dumb, but nearly two years on, the benefits to Google of its great AI giveaway are increasingly evident. Today TensorFlow is becoming the clear leader among programmers building new things with machine learning. "We have significant usage today, and it's accelerating," says Jeff Dean, who led TensorFlow's design and heads Google's core artificial-intelligence research group. Once you've built something with TensorFlow, you can run it anywhere—but it's especially easy to transfer it to Google's cloud platform. The software's popularity is helping Google fight for a bigger share of the roughly $40 billion (and growing) cloud infrastructure market, where the com-

5. Alp



LEONARD GRECO

pany lies a distant third behind Amazon and Microsoft.

The head of Google's cloud business, Diane Greene, said in April that she expects to take the top spot within five years, and a core part of Google's strategy for catching up is to appeal to the sudden enthusiasm about artificial intelligence in industries from health care to autos. Companies investing in the technology are expected to spend heavily with cloud providers to avoid the costs and complexity of building and running AI themselves, just as they pay today for cloud hosting of e-mail and websites. Customers like insurer AXA—which used TensorFlow to make a system that predicts expensive traffic accidents—also get the benefits of the same infrastructure Google uses to power their own products. Google says that means better performance at competitive prices. S. Somasegar, a managing director at venture fund Madrona who was previously head of Microsoft's developer division, says TensorFlow's prominence poses a genuine challenge to Google's cloud rivals. "It's a fantastic strategy—Google is so far behind in cloud, but they've picked an area where they can create a beachhead," he says.

Inside Google, TensorFlow powers products such as the Google Translate mobile app, which can translate a foreign menu in front of your eyes when you point your phone at it. The company has created specialized processors to make TensorFlow faster and reduce the power it consumes inside Google's data centers. These processors propelled the historic victory of software called AlphaGo over a champion of the ancient board game Go last year and are credited with making possible a recent upgrade that brought Google's translation service close to human level for some languages.

TensorFlow is far from the only tool out there for building machine-learning software, and experts can argue for hours about their individual merits. But the weight of Google's brand and its technical advantages make its package stand out, says Reza Zadeh, an adjunct professor at Stanford. He originally built his startup Matroid, which helps companies create image recognition software, around a competing tool called Caffe, but he dumped it after trying TensorFlow. "I saw it was very clearly superior in all the technical aspects, and we decided to rip everything out," he says.

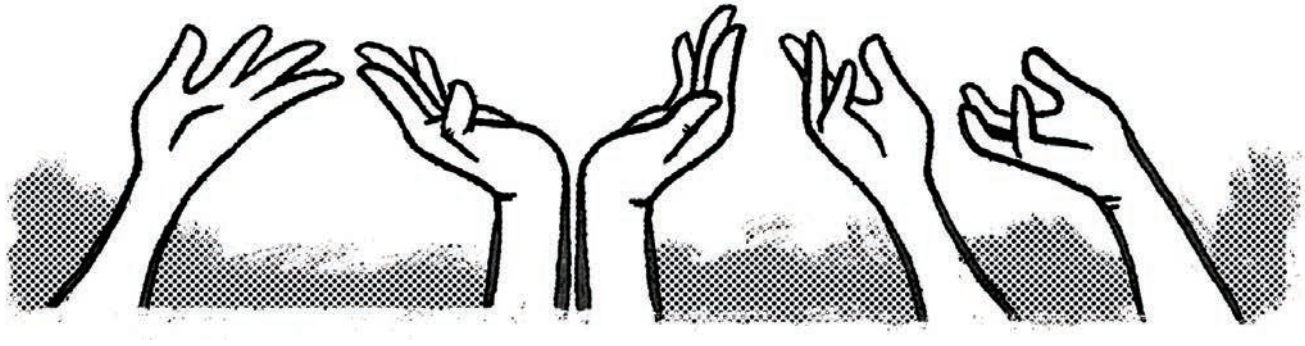
Google's tool is also becoming firmly lodged in the minds of the next generation of artificial-intelligence researchers and entrepreneurs. At the University of Toronto, an AI center that has schooled many of today's leading researchers, lecturer Michael Guerzhoy teaches TensorFlow in the university's massively oversubscribed introductory machine-learning course. "Ten years ago, it took me months to do something that for my students takes a few days with TensorFlow," says Guerzhoy.

Since Google released TensorFlow, its competitors in cloud computing, Microsoft and Amazon, have released or started supporting their own free software tools to help coders build machine-learning systems. So far, says Guerzhoy, neither has as broad and dedicated a user base as TensorFlow among researchers, students, and working coders.



ADAM NICKEL





18. Carbon

CEO Joseph DeSimone has taken a whole new approach to advanced manufacturing, and it's working.

By Katherine Bourzac

A sleek mechanical arm plunges into a pool of what looks like milky gray ink in Carbon's lab in Redwood City, California. The black arm slowly moves upwards, pulling a latticed plastic cube out of the bath, shiny and dripping with ink: a large-scale model of the porous structure of bone.

Joseph DeSimone, Carbon's CEO and cofounder, looks on. DeSimone, a polymer chemist, helped invent these machines, and he still gets a kick out of watching them work. It is a form of 3-D printing, but it's done in a novel way that is faster than previous techniques and works with many more types of plastics. The way the printer pulls the object smoothly from the pool of milky liquid, millimeter by millimeter, gives the illusion that an existing structure is emerging. In fact, the liquid is a light-sensitive precursor material; a digital projector is continuously projecting ultraviolet light onto the bottom of the lattice, the first of two steps that harden the material to form the plastic object.

Using its process for rapidly printing objects with high-performance polymers like polyurethanes and epoxies, four-year-old Carbon is pursuing an approach fundamentally different from other methods of 3-D printing, which put down layers of plastic one at a time. That limits the quality of many 3-D-printed products, but Carbon fabricates an object in a continuous process, avoiding some typical problems. DeSimone say his technology enables Carbon to print polymer objects rapidly, in some cases thousands of times faster than other 3-D printers, and use a wider range of materials, including rubber-like elastomers and durable, hard plastics.

Technology for 3-D printing has been around since the 1980s; fans have been intrigued by the potential to easily generate objects with complex structures that can't be made using other techniques, and to personalize objects like medical devices or apparel as easily as different images can be printed on a piece of paper. However, two challenges have kept 3-D printing a niche technology: its slow speed and the limited range of materials for which it's been suited. Recently, companies like HP, which sells 3-D printers to fabricate high-strength nylon parts, and Desktop Metal (see "The 3-D Printer That Could Finally Change Manufacturing," May/June), whose machine uses various metal alloys, have attempted to make the technology more relevant for manufacturing by increasing its speed and using industrially desirable materials. Carbon joins this increasingly crowded field. Backed with $222 million from investors including GE Ventures and Alphabet's Google Ventures, the company has created a process it calls "digital light

18. Ca



LEONARD GRECO

synthesis" rather than 3-D printing, a way of alluding to its novel chemistry.

Adding credibility to DeSimone's claim that 3-D printing is finally ready to mass-manufacture customized parts is Carbon's growing list of paying customers. "The volume and speed we can achieve with Carbon's pioneering digital light synthesis is unprecedented," says Paul Gaudio, the global creative director of Adidas. "It's almost magical." Adidas tried other 3-D printers as it sought to make customized shoes for the mass market, but none of them could operate at large scale. Now it will use Carbon's technology to manufacture elastomer midsoles for athletic shoes, starting with 5,000 pairs this fall and winter. The company hopes to print millions of shoes, including custom-tailored models, by 2021. Other customers are taking advantage of Carbon's technology to print parts for electric motorcycles, server farms, and cooling systems, all of which have been difficult or impossible to make with other methods.

Blending Polymers

Before he founded Carbon in 2013, DeSimone spent more than 20 years as a polymer chemist at the University of North Carolina. In 2012, as 3-D printing was becoming popular with hackers using desktop models, he received a call from Alex Ermoshkin, a former employee of his lab and of a drug delivery company DeSimone had started up in the early 2000s. Ermoshkin came to him with a business idea: would he be interested in starting a company to build cheaper 3-D printers?

DeSimone, who had some experience with the technology, was intrigued. He suggested that Ermoshkin do a patent search to see what was missing in the field. Ermoshkin found many descriptions of ways to print three-dimensional objects one layer at a time, some using UV light to link up chemical building blocks and harden them to form polymers. Stereolithography is one of the original 3-D-printing methods, but it has always been limited by the properties of such UV-curable plastics, which have not included durability or flexibility. Also, resulting products formed in layers tended to be weak. Stereolithography is useful for making prototypes, but not for final parts. What's more, the printers were very slow, taking hours or even days to make large objects.

The two began to think about how to print without stopping. They were inspired, DeSimone says, by the image of the T-1000 humanoid robot rising up out of a pool of metallic liquid to stand on two feet in the 1991 movie *Terminator 2: Judgment Day*. In the system that DeSimone and Ermoshkin imagined, light patterns would be projected into the bath, like images on a movie screen, as the printer arm moved up. As the light hardened the material, layers would form continuously.

Among the tricky challenges the team had to solve was how to keep objects from sticking to the bottom of the bath. DeSimone came up with the solution: an oxygen-permeable window, which takes advantage of the fact that these types of UV-triggered reactions are stopped by oxygen. It's a small detail, but a critical one.



Carbon's vice president of materials, Jason Rolland, says the company has focused on overcoming the limitations of UV-curable polymers, developing blends that offer a wide range of properties. The system can now print with 12 classes of materials—some durable, some stretchy or squishy, some able to bear heavy weights. One, a cyanate ester that can withstand temperatures up to 220 °C, is suitable for making automotive and aerospace parts. Another class, elastomers, includes the squishy materials Adidas is using. DeSimone foresees "an app store for resins" from which the company's customers can order to print what they need.

Scaling Up

Carbon worked with Adidas on 150 different iterations of the elastomer in its new shoe, which uses a printed midsole with a lattice structure. The structure's mechanical properties can be tailored by changing the pattern of the lattice's holes and struts, bringing personalization to a mass scale.

In Carbon's lobby, a large display screen shows the operating status of all its printers, both on site and with customers. The company sells the devices on a subscription model and then works with customers to select the right materials and designs. Every six weeks or so, Carbon sends out software updates based on what it is seeing in the field—a tactic inspired by Tesla, the former employer of Carbon's VP of engineering, Craig Carlson. Carbon is now designing machine-learning software to help generate the optimal printing design and parameters for any given product. "We want to see the problems before we print," DeSimone says.

ADAM NICKEL; JOSHUA MATHEWS



The Adidas Futurecraft 4D includes midsoles printed by Carbon.

32. Oxford Nanopore

A British company bets a hand-held analyzer will change how we look at DNA.

By Antonio Regalado

Last June, a British-Brazilian scien-tific team boarded a bus that had been turned into a makeshift laboratory and headed out to tour six cities across northeastern Brazil.

The researchers were there to find mosquitoes infected with the Zika virus and sequence its genome in their blood, since the evolution of the viral genome contains clues to the epidemic's origins. But rather than collect insects and send the samples back to a central lab, they'd outfitted the bus with everything they needed to do the research. The most important item: a DNA sequencer about the size and weight of a deck of cards that runs off a laptop USB plug and costs just $1,000.

The instrument, called the MinION, works by pulling DNA through around 500 nanoscopic pores and reading it as it passes through by measuring an electrical signal produced by each nucleotide, or DNA letter. Oxford Nanopore Technologies,

32 Ox Na



LEONARD GRECO

a private British company that has spent 12 years and about $200 million developing the invention, foresees its cheap DNA sequencers providing a way to study life in real time. Every living thing has its own distinctive DNA, and being able to read it provides a powerful tool to identify and profile microbes in great detail. In the last year, the same gadget has seen action in Antarctica to check for life in frozen valleys, on the International Space Station to sequence genes in space for the first time, and deep below the earth in a Welsh coal mine called the "Big Pit."

A central question for Oxford as a business is whether such portable sequencing is widely useful. Privately valued at around $1.5 billion, Oxford introduced the MinION in 2014, but its sales remain negligible, at around $6 million last year, and the company sometimes gives away the supplies needed to run the instrument. Still, it is betting that users of the device will create new applications, thereby expanding demand for the technology and allowing the company to conquer the broader sequencing market.

Globally, the market for high-speed DNA sequencers and the chemicals to run them is currently about $3 billion a year, according to the consultancy DeciBio—roughly the amount that a single blockbuster drug brings in. The market is dominated by Illumina, a San Diego company whose top-of-the-line machine costs $1 million, takes up as much space as a large filing cabinet, and weighs 498 pounds. It can decode 35 human genomes per week, at high accuracy, for less than $1,000 apiece. In the hands of big academic centers and companies, these machines are powering genetic research into the causes of disease and, increasingly, the quest for new types of cancer diagnostics and prenatal tests.

All those activities might one day be run instead using nanopore sequencing—or at least Oxford hopes so. Right now, the MinION is a good way to identify and study bacteria and viruses. But last year, the first human genomes were sequenced on a MinION as well. (It's still a cumbersome process and not as inexpensive, however—taking close to $20,000 worth of disposable cartridges at $500 each.)

Three years ago, when Oxford began sending out the first MinIONs to select labs to try out, the devices' performance was shaky at best. Early models were error prone and often didn't work at all. But these early mistakes, Oxford says, were the price of a strategy to recruit tech-savvy biologists who have helped it improve the device, create buzz, and explore potential applications. And Oxford has been improving the MinION rapidly—in two years, the device has become far more accurate and five times as fast. The DNA strands now barrel through each pore at 450 letters per second. And there's headroom for further improvements. "I can see it happening. It's developing by leaps and bounds," says Benedict Paten, a bioinformatics expert at the University of Southern California. "I really believe they are about to dominate sequencing."

For now, the MinION isn't as fast or accurate as Illumina's machines. But sequencing typically involves observing DNA with big, costly microscopes and elaborate chemistry. In nanopore sequencing, by contrast, DNA strands are sucked through minuscule ring-shaped biological pores borrowed from the surface of bacteria. As the strings of DNA letters—known as A, G, C, and T—pass through, each generates a faintly shifting electrical signal that's usually strong enough to let the machine identify it.



Top: Jaqueline Goes de Jesus of the Oswaldo Cruz Foundation and Nuno Faria of the University of Oxford use the MinION sequencer in João Pessoa, Brazil, as part of a mobile study on the origins of the Zika virus.

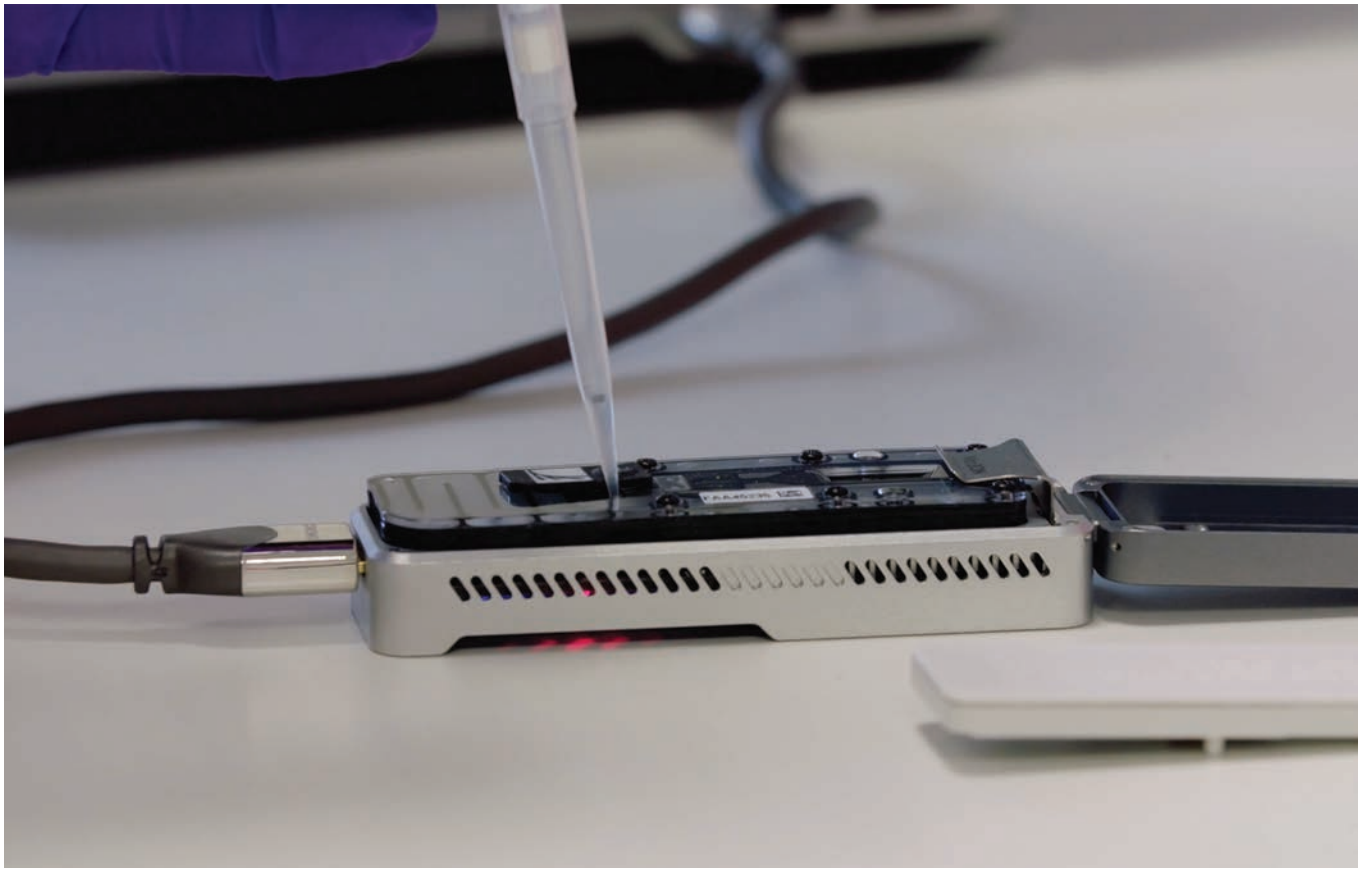
Bottom: Bruna do Nascimento of the Evandro Chagas Institute in Brazil studies Aedes aegypti *mosquitoes that spread Zika.*

ADAM NICKEL





COURTESY OF ZIBRA PROJECT





Above: The MinION DNA sequencer weighs 100 grams and is powered by a USB cord.

Left: Astronaut Kate Rubins sequences DNA aboard the International Space Station in 2016.

COURTESY OF NASA PUBLIC AFFAIRS; COURTESY OF OXFORD NANOPORE

This April, Oxford flew many of its key users to London for its annual event, a kind of biological developers' conference. At the conference, chief technology officer Clive Brown took the stage to alternately boast about and apologize for a parade of half-finished ideas for making the MinION better and creating future products.

Starting with what Brown now calls the "notorious" and premature announcement of the MinION itself back in 2012 (it didn't appear for another year and a half), Oxford has frequently made big promises and then missed deadlines. Twice since 2016 it dodged patent infringement lawsuits brought by Illumina and another competitor, Pacific Biosciences, by changing key aspects of the MinION. In both cases, the updates ended up making the device better, not worse.

"There seems to be a halo around this company," says its patent lawyer, Martyn Andrews.

MinION can do some unusual tricks. One is to read out extraordinarily long stretches of consecutive DNA letters. Illumina's instruments read DNA in short fragments of 150 letters. With a MinION, readouts of 10,000 letters are common. Its record: 882,000 DNA letters in a row.

Being able to read long sequences of letters eases the job of digitally reassembling a genome from its parts: imagine a puzzle with thousands of pieces to fit together rather than millions. That's a crucial factor when exploring an organism's genome for the first time. It also makes it simpler to study patients' DNA for certain disorders and cancers caused when genes get duplicated or deleted—what scientists call "structural variation," as opposed to mutations that affect single letters.



ADAM NICKEL

While MinION is mostly suitable for analyzing bacteria and viruses, it's a second device Oxford is developing, the Promethion, that Brown believes will be Oxford's "Illumina killer," because it will be suitable for sequencing humans and other organisms with large genomes. A printer-size instrument that can route DNA through tens of thousands of pores at once, the Promethion will read a trillion DNA letters in a few hours, Brown claims—similar to Illumina's high-end $1 million sequencing instrument, called the HiSeq X—but won't cost nearly as much to buy.

Like many of Oxford's ideas, Promethion is a work in progress. And it's already well behind schedule. Oxford now plans to start selling it later this year. "Just sit tight," Brown told scientists at Oxford's meeting. "The raw horsepower is enough to put you ahead."

One team in possession of an early model is ZF-Screens, a company in the Netherlands, where scientists hope to decode the genome of the country's national flower, the tulip. Knowing this commercially valuable flower's genetic details could speed up the creation of new varieties, which can take as long as 30 years because the tulip's maturation time is so long. But the plant's genome is huge—about 10 times the size of a human's—and so highly repetitive that it can't easily be decoded using existing sequencing methods. Team member Hans Jansen says he has tested the Promethion and "it kind of works and kind of doesn't work." The software to capture data is poorly worked out, and the first batch of pore cells arrived broken.

But Jansen plans to stick with it. That's because he thinks the longer DNA readouts may be the only way to conquer the tulip's labyrinthine genome. "We need this technology," he says.

39. IBM

Although it has overhyped its Watson machine-learning system, the company still could have the best access to the kind of data needed to make medicine much smarter.

By David H. Freedman

Paul Tang was with his wife in the hospital just after her knee replacement surgery, a procedure performed on about 700,000 people in the U.S. every year. The surgeon came by, and Tang, who is himself a primary-care physician, asked when he expected her to be back at her normal routines, given his experience with patients like her. The surgeon kept giving vague non-answers. "Finally it hit me," says Tang. "He didn't know." Tang would soon learn that most physicians don't know how their patients do in the ordinary measures of life back at home and at work—the measures that most matter to patients.

Tang still sees patients as a physician, but he's also chief health transformation officer for IBM's Watson Health. That's the business group developing health-care applications for Watson, the machine-learning system that IBM is essentially betting its future on. Watson could deliver information that physicians are not getting now, says Tang. It could tell a doctor, for instance, how long it took for patients similar to Tang's wife to be walking without pain, or climbing stairs. It could even help analyze images and tissue samples and determine the best treatments for any given patient.

It's because of possibilities like this that health care is one of the hottest segments of the market for machine-learning technologies. Research firm CB Insights counts at least 106 startups that have sprung up since 2013 and are still in business.

None of those companies has garnered anywhere near the attention that Watson has, thanks to its victory on the television quiz show *Jeopardy!* in 2011 and zealous marketing by IBM ever since. But lately, much of the press for Watson has been bad. A heavily promoted collaboration with the M.D. Anderson Cancer Center in Houston fell apart this year. As IBM's revenue has swooned and its stock price has seesawed, analysts have been questioning when Watson will actually deliver much value. "Watson is a joke," Chamath Palihapitiya, an influential tech investor who founded the VC firm Social Capital, said on CNBC in May.

However, most of the criticism of Watson, even from M.D. Anderson, doesn't seem rooted in any particular flaw in the technology. Instead, it's a reaction to IBM's overly optimistic claims of how far along Watson would be by now. In fact, it still

MIT TECHNOLOGY REVIEW
VOL. 120 | NO. 4

39

BM



LEONARD GRECO

It seems likely that Watson Health will be a leader in applying AI to health care's woes. If Watson has not, as of yet, accomplished a great deal along those lines, one big reason is that it needs certain types of data to be "trained." And in many cases such data is in very short supply or difficult to access. That's not a problem unique to Watson. It's a catch-22 facing the entire field of machine learning for health care.

Though the problem of missing and inaccessible data may slow Watson down, it may hurt IBM's competitors more. That's because the best bet for getting the data lies in close partnerships with large health-care organizations that tend to be technologically conservative. And one thing IBM still does very well in comparison to start-ups, or even giant rivals like Apple and Google, is gain the trust of executives and IT managers at big organizations. The specific problems with the M.D. Anderson project notwithstanding, IBM has a crucial advantage. It's getting Watson inside a wide range of medical centers, health-care administration groups, and life-science companies, all of which are positioned to provide the critical data needed to shape AI's future in medicine.

"Health care has been an embarrassingly late adopter of technology," says Manish Kohli, a physician and health-care informatics expert with the Cleveland Clinic.

Unrealistic timelines

The breakup with M.D. Anderson seemed to show IBM choking on its own hype about Watson.

The cancer center and IBM partnered in 2012. The goal was for Watson to read data about any patient's symptoms, gene sequence, and pathology reports, combine it with physicians' notes on the patient and relevant journal articles, and then help doctors come up with diagnoses and treatments. But IBM and M.D. Anderson both overinflated expectations for the technology. IBM claimed in 2013 that "a new era of computing has emerged" and gave *Forbes* the impression that Watson "now tackles clinical trials" and would be in use with patients in just a matter of months. In 2015, the *Washington Post* quoted an IBM Watson manager describing how Watson was busy establishing a "collective intelligence model between machine and man." The *Post* said that the computer system was "training alongside doctors to do what they can't."

In February of this year, the University of Texas, which runs M.D. Anderson, announced it had shuttered the project, leaving the medical center out $39 million in payments to IBM—for a project originally contracted at $2.4 million. After four years it had not produced a tool for use with patients that was ready to go beyond pilot tests. M.D. Anderson wouldn't comment to me about Watson specifically, but it appears that the problems stemmed mainly from internal struggles over how the project was managed and funded.

That's not to say IBM has no troubles with Watson. Indeed, they're larger than what any one implementation reveals.

To understand what's slowing the progress, you have to understand how machine-learning systems like Watson are trained. Watson "learns" by continually rejiggering its internal processing routines in order to produce the highest possible percentage of correct answers on some set of problems, such as which radiological images reveal cancer. The correct answers have to be already known, so that the



ADAM NICKEL

system can be told when it gets something right and when it gets something wrong. The more training problems the system can chew through, the better its hit rate gets.

That's relatively simple when it comes to training the system to identify malignancies in x-rays. But for potentially groundbreaking puzzles that go well beyond what humans already do, like detecting the relationships between gene variations and disease, Watson has a chicken-and-egg problem: how does it train on data that no experts have already sifted through and properly organized? "If you're teaching a self-driving car, anyone can label a tree or a sign so the system can learn to recognize it," says Thomas Fuchs, a computational pathologist at Memorial Sloan-Kettering, a cancer center in New York. "But in a specialized domain in medicine, you might need experts trained for decades to properly label the information you feed to the computer."

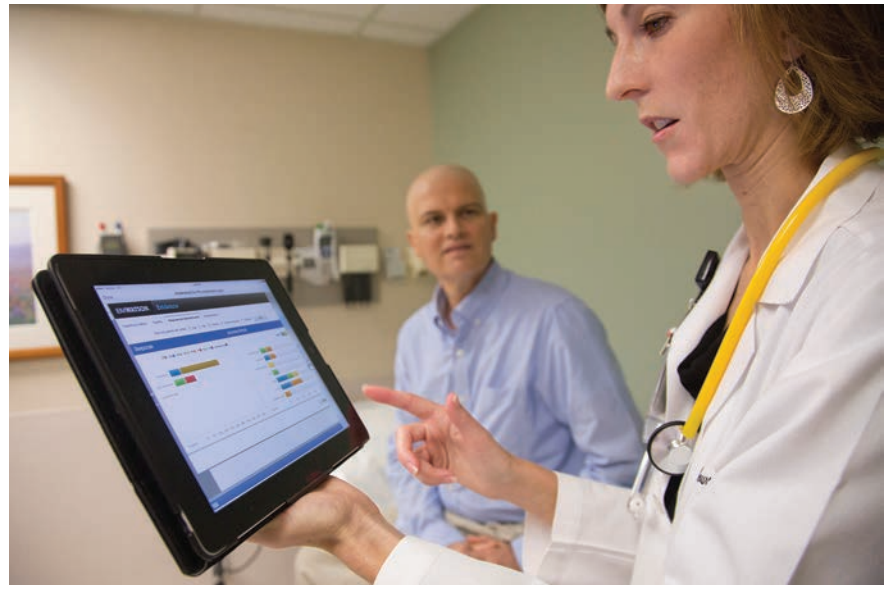
Some version of that stumbling block emerges in every domain in which IBM hopes to have Watson contribute—as it does for any company's machine-learning solution. To train Watson to go through giant pools of data and pull out the few pieces of information important to a single patient, someone has to do it by hand first, for thousands and thousands of cases. To recognize genes linked to disease, Watson needs thousands of records of patients who have specific diseases and whose DNA has been analyzed. But those gene-and-patient-record combinations can be hard to come by. In many cases, the data simply doesn't exist in the right format—or in any form at all. Or the data may be scattered throughout dozens of different systems, and difficult to work with.

Consider, for example, the goal of improving primary care by placing better data at the fingertips of clinicians. When doctors miss chances to treat relatively minor concerns during a routine primary-care visit, before a more advanced problem sends patients to an emergency room or a specialist, their health suffers and costs explode. "About one-third of every dollar spent on health is probably unnecessary," says Anil Jain, IBM Watson Health's chief medical officer, who is also a practicing primary-care physician. Machine learning is widely recognized as an opportunity to address that problem.

To really help doctors get better outcomes for patients, however, Watson will need to find correlations between what it reads in health records and what Tang calls "all the social determinants of health." Those factors include whether patients are drug-free, avoiding the wrong foods, breathing clean air, and on and on. But Tang concedes that today almost no hospitals or medical practices get that data reliably for a significant percentage of patients. Part of the problem is that hospitals have been slow to take up modern, data-driven practices. "Health care has been an embarrassingly late adopter of technology," says Manish Kohli, a physician and health-care informatics expert with the Cleveland Clinic.

Where the data does exist, IBM has often simply gone out and bought it. It has acquired companies such as Truven Health Analytics, Explorys, and Phytel, all of which were active in dealing with large data sets across hospitals and patient populations. And even after the demise of the M.D. Anderson deal, IBM has some critical partnerships that further its access to patient data.

One of them is with Atrius Health, a network of nearly 900 mostly primary-care physicians throughout the Boston area. The partnership aims to develop and test a Watson-based system capable of pulling out nuggets of information critical to an individual patient from an ocean of notes, records, and articles. "Seeking all the relevant information is an onerous job for primary-care physicians as things exist today," says Atrius's chief medical officer, Joe Kimura. Electronic health records may have made the problem even worse, he adds, because the advent of such systems has enormously increased the amount of data generated in each visit, without providing a standard format to allow easy retrieval.



A leukemia doctor at M.D. Anderson, Courtney DiNardo, used IBM's Watson system while consulting with a patient in 2013.

Critically, many of the most important notes in patient records are sentences that a conventional IT system can't make sense of. But Watson can apply the natural-language-processing skills developed for *Jeopardy!* in order to extract meaning from them. Ideally, it could then suggest ways physicians can help patients avoid the need for extensive care. "Why should we focus only on making sure we've done as good a job as possible with patients who break a hip," asks Kimura, "when we can try to predict which patients are at risk for falls and help them not break the hip at all? We need to push our care upstream."

Watson Health is also partnering with the Central New York Care Collaborative, a state-government-funded agency that works with some 2,000 health-care providers in six counties. The partnership is intended to support the goal of a 25 percent reduction in emergency room admissions and hospital readmissions, when patients who have been discharged from a hospital have to return to address related problems. It also provides potential access to vast amounts of patient data.

There are other ways to get such data. One of Google's sister companies is trying to get it directly from patients themselves. Verily Life Sciences, a health-care divi-



ADAM NICKEL; THOMAS SHEA/IBM

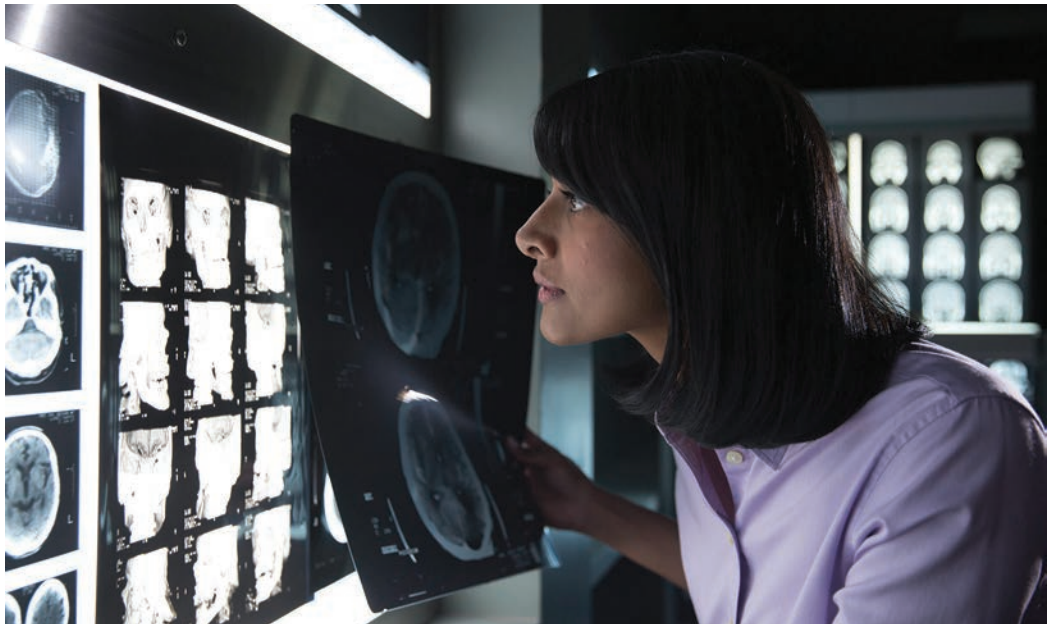
sion of Alphabet, is partnering with Duke and Stanford to develop a highly structured health database on some 10,000 volunteers. The database will be filled with information not only from their clinical visits but also from wearable health-monitoring devices. That could be a promising leap in data access, though it could take a decade or more to produce highly usable results.

Fuchs's group at Memorial Sloan-Kettering hopes to train an AI system to read tissue-stain slides, a process that will require a large library of digitally annotated slides with confirmed diagnoses and other critical data. So the group is gearing up to produce 40,000 such slides a month on its own. "That's far more than anyone else," Fuchs says. "It's an enormous task because of all the variability in biology."

Even M.D. Anderson, despite the fate of the Watson project, is continuing a large program that began around the same time, focused on gathering 1,700 types of clinical data on every patient who walks in its doors. Andy Futreal, the scientist who runs the program, says combining that patient information with research data will be crucial for the sorts of capabilities that systems like Watson could provide. "Once we have the data in place, now you can get into the business of AI machine learning uncovering those factors dictating who does and who doesn't do well with different treatments," says Futreal.

IBM, for its part, continues to rack up data from partnerships. For cancer diagnosis and care alone, the company has partnered with Memorial Sloan-Kettering, the Mayo Clinic, the Harvard- and MIT-affiliated Broad Institute, and medical-test giant Quest Diagnostics. The Memorial Sloan-Kettering collaboration has already produced a system that sifts through journal literature to inform treatment decisions, and it has been rolled out to the Jupiter Medical Center in Florida and a hospital chain in India. On the drug-discovery front, Watson Health is working with the Barrow Neurological Institute, where Watson helped find five genes linked to ALS that had never before been associated with the disease, and with the Ontario Brain Institute, where Watson identified 21 promising potential drug candidates.

Will Watson eventually make a difference in improving health outcomes and lowering costs? Probably, says Stephen Kraus, a partner at the VC firm Bessemer Venture Partners who focuses on health care and has invested in AI health-care startups. "It's all for real," says Kraus. "This isn't about putting out vaporware in order to boost stock prices." But Kraus joins most experts in cautioning against unrealistic timelines or promises—some of which have come from IBM itself. "This is hard," he says. "It's not happening today, and it might not be happening in five years. And it's not going to replace doctors."



COURTESY OF IBM

IBM announced in 2015 that Watson's diagnostic capabilities would be boosted by data obtained from Merge Healthcare, a medical imaging management company that IBM bought for about $1 billion.

40. General Electric

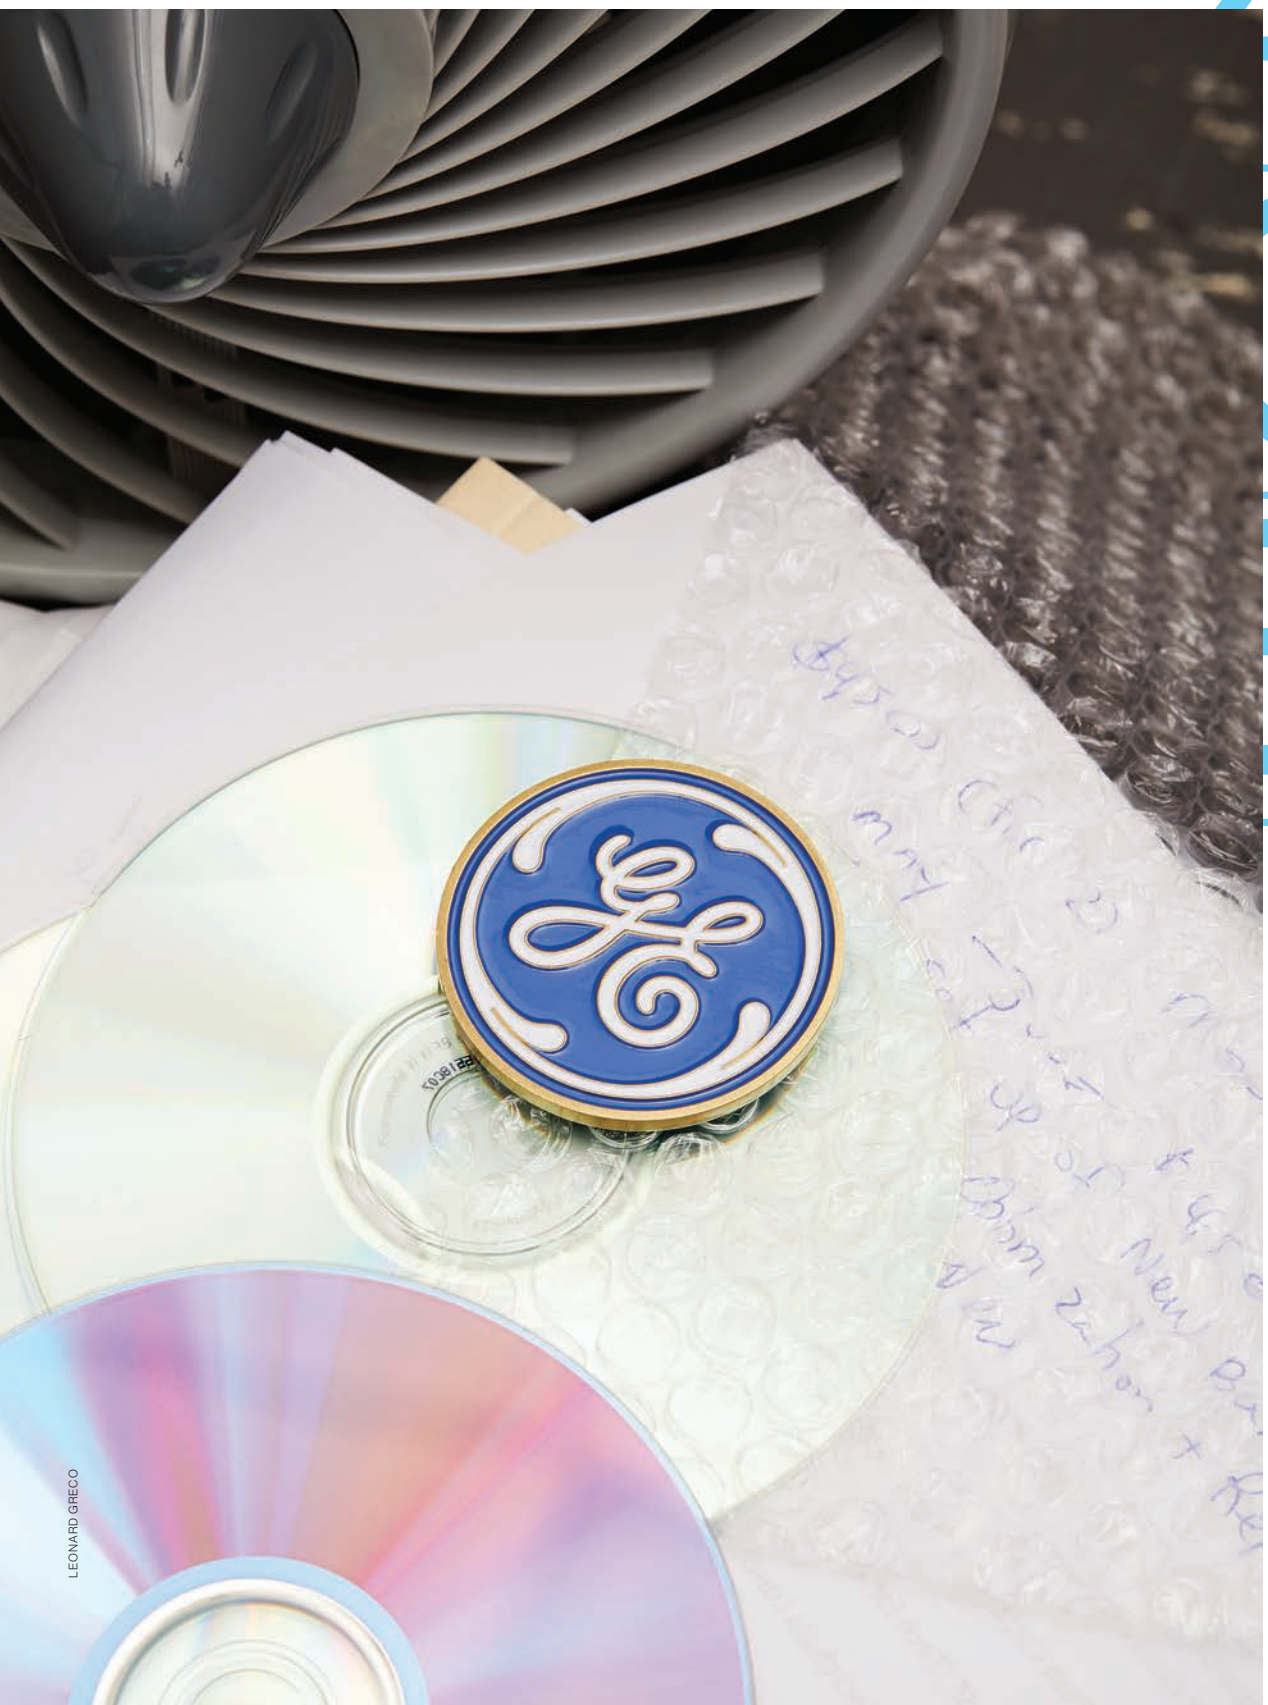
As part of its shift toward high-tech businesses, the 125-year-old company is threading artificial intelligence throughout its operations, starting with its scientists.

By Elizabeth Woyke

When Jason Nichols joined GE Global Research in 2011, soon after completing postdoctoral work in organic chemistry at the University of California, Berkeley, he anticipated a long career in chemical research. But after four years creating materials and systems to treat industrial wastewater, Nichols moved to the company's machine-learning lab. This year he began working with augmented reality. Part chemist, part data scientist, Nichols is now exactly the type of hybrid employee crucial to the future of a company working to inject artificial intelligence into its machines and industrial processes.

Fifteen years ago, GE's machine operators and technicians monitored its aircraft engines, locomotives, and gas turbines by listening to their clanks and whirs and checking their gauges. Today, the company uses AI to do the equivalent, even predicting failures in advance. By marshaling this technology, GE hopes to become one of the world's top software providers by 2020, a quest that amped up in 2011 with a $1 billion initiative to collect and analyze sensor data from machines. Creating smarter models via AI is the next step in the company's strategy – one that it hopes will give it an advantage over longtime rivals like Siemens and soft-

40 GE



LEONARD GRECO

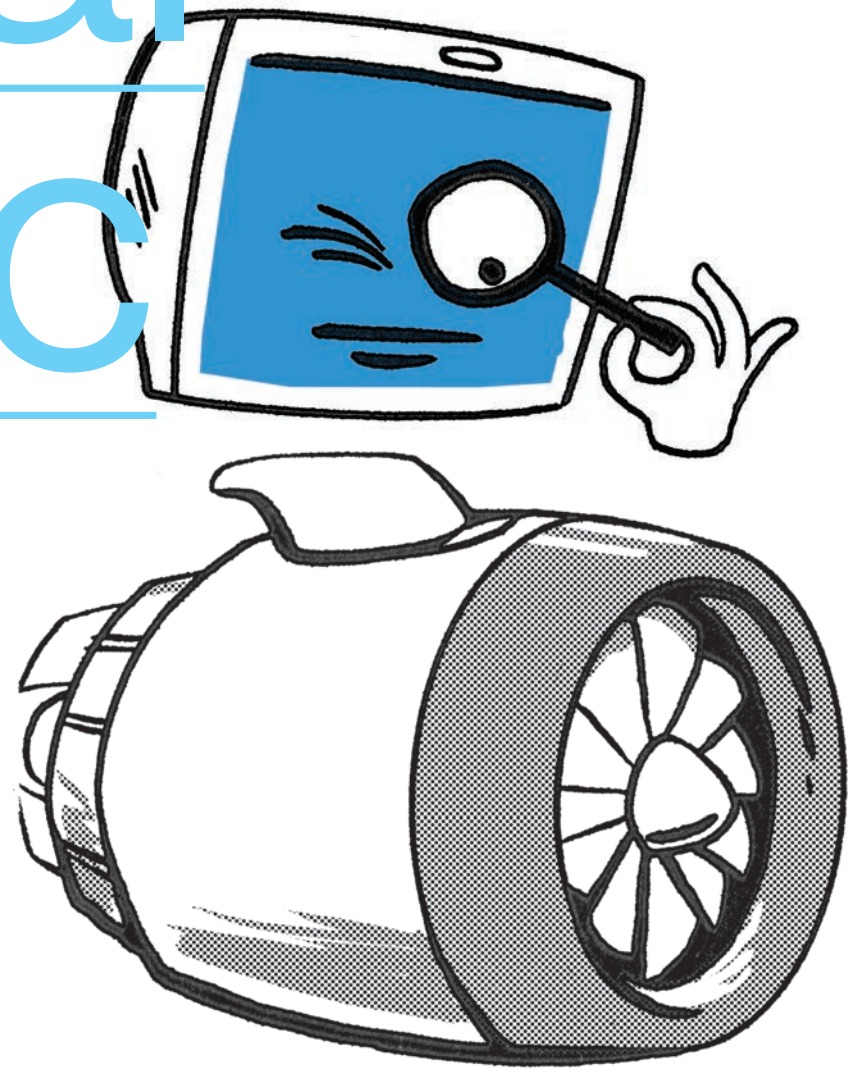
ware giants, such as IBM, that are now expanding into industrial analytics.

Of course, integrating artificial intelligence into an organization founded in 1892 is a difficult task. It starts with training the technical brains behind the company, which employs 300,000 people across all its businesses worldwide. GE Global Research, where Jason Nichols works, is setting up online programs that teach machine learning and symposia where scientists can explore new roles. So far, nearly 400 employees from across the company have completed GE's certification program for data analytics, and about 50 scientists have moved into digital analytics jobs of the kind Nichols has taken on.

Dual careers

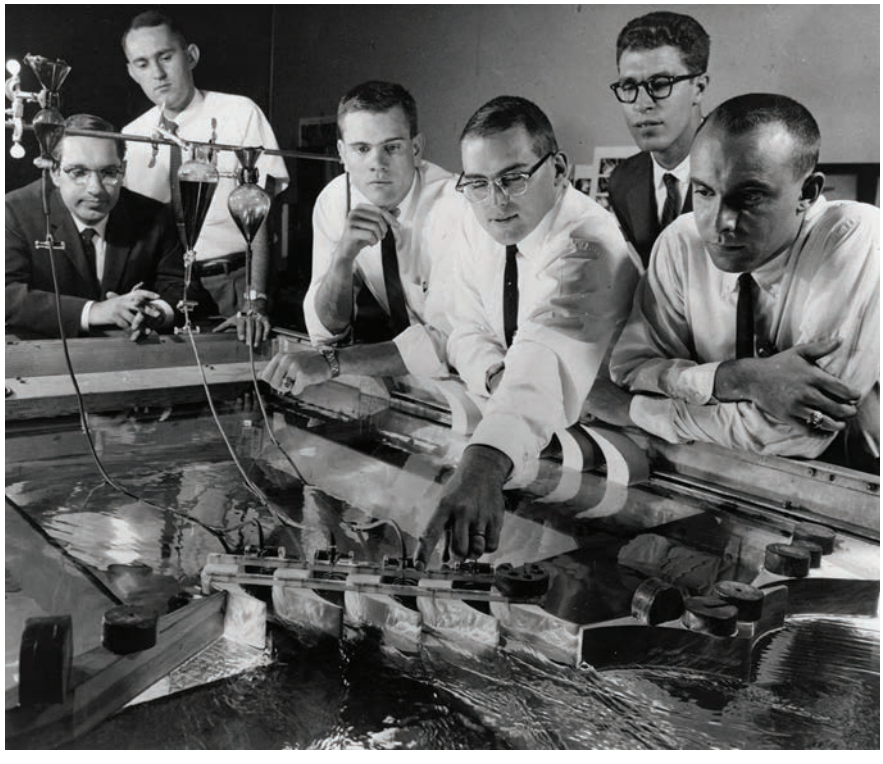
Many of these dual scientists help make cloud-hosted software models of GE's machines that can be used to save money and improve safety for its customers. GE builds these "digital twins" using information it gathers from sensors on the machines, supplemented with physics-based models, AI, data analytics, and knowledge from its scientists and engineers. Though digital twins are primarily lines of software code, the most elaborate versions look like 3-D computer-aided design drawings full of interactive charts, diagrams, and data points. They enable GE to track wear and tear on its aircraft engines, locomotives, gas turbines, and wind turbines using sensor data instead of assumptions or estimates, making it easier to predict when they will need maintenance. An aircraft engine flying over the U.S. could, for instance, have a digital twin on a GE computer server in California help determine the best service schedule for its parts.

Besides forecasting a machine's life expectancy, the virtual models allow GE



ADAM NICKEL





FOX PHOTOS/GETTY

to optimize the operation of its products. GE says digital twins are increasing the amount of electricity wind farms produce by as much as 20 percent and reducing annual fuel consumption and carbon emissions for one of its locomotives by 32,000 gallons and 174,000 tons a year, respectively. More than 700,000 models have been delivered to clients, a number that could exceed one million by the end of this year.

The technology depends on artificial intelligence to continually update itself. What's more, if data is corrupted or missing, the company fills in the gaps with the aid of machine learning, a type of AI that lets computers learn without being explicitly programmed, says Colin Parris, GE Global Research's vice president for software research. Parris says GE pairs computer vision with deep learning, a type of AI particularly adept at recognizing patterns, and reinforcement learning, another recent advance in AI that enables machines to optimize operations, to enable cameras to find minute cracks on metal turbine blades even when they are dirty and dusty.

Top: GE scientists have been studying and adapting changing technology for decades. In this photo from April 18, 1968, a GE systems engineer mans the telex link of a Student Response System at Syracuse University. The computer analyzed the answers given by students to multiple-choice questions and relayed them back to the teacher.

And on March 20, 1969, GE researchers studied the flow of gases, invisible in their natural state, by creating simulations on a device called a working water table, where currents were observed through the use of dyes.



GE uses AI to create continuously updating digital representations of its machines, such as this gas turbine, the 9HA, at a plant in Belfort, France.

FREDERICK FLORIN/GETTY

Take the tiny robot, a little bigger than a Matchbox car, used to inspect working engines. Using computer vision and a variety of AI techniques, the bot can look for cracks inside plane engines by riding on top of a slowly moving fan blade.

Similar technology can be attached to a drone to find corrosion on the 200-foot-high flare stacks that burn off excess gas released at oil and gas production sites.

No fad

To develop and work with these systems, GE researchers need to understand both the physics of the machines and the AI algorithms.

"This is a place where you will have a molecular biologist sitting with a machine-learning expert or a controls-theory person sitting with someone who knows about materials science," says Mark Grabb, GE Global Research's technology director for analytics. "That type of collaboration is very powerful, but there is nothing more powerful than having that same information in the same brain; it's just hyper-efficient."

Consider the brain of Matt Nielsen, who joined GE Global Research in 1998 after earning a PhD in physics. Nielsen developed photonics and worked on electric-vehicle software before moving fully to the company's digital side in 2015. Today, he



Digital replicas of jet engines help GE's aviation customers save money by predicting exactly when they will need maintenance. Here a GE engine sits at an overhaul facility in Rio de Janeiro, Brazil.

YASUYOSHI CHIBA/GETTY

leads a team of digital twin developers and helps build physics-based models that can be combined with machine-learning algorithms.

Sahika Genc, another dual scientist, developed systems for ICU alarms before transitioning to GE's machine-learning lab in 2014. Genc is now a machine-learning scientist who uses deep learning and reinforcement learning to make GE's energy management systems more efficient. One of her recent projects applied machine learning and heat transfer theory to identify how building energy is dissipated and stored. The forecasts will help GE customers reduce their energy consumption.

These hybrid researchers could be GE's best shot at remaining relevant for another century as the company looks for growth opportunities in such competitive and mature industries as turbines, jet engines, and locomotives.

Parris, the software research leader, admits that some of GE's 2,000 researchers still regard certain aspects of the new approach as a "passing fad."

But scientists who don't make the leap may get left behind. In January, the company laid off researchers in areas deemed peripheral to GE's "digital industrial" strategy. That's after creating 100 new research jobs related to AI and robotics in 2016.



DAVID BRANDON GEETING

Reviews

Virtual Reality's Missing Element: Other People

VR can be the basis of a new communications industry if the technology becomes less insular and isolating.

By Rachel Metz

Unlike most people, I have a virtual-reality headset. I have exactly one friend who also has one. So most of the time I spend in VR, I'm all by myself. I can almost hear the digital tumbleweeds rolling by.

That's a funny thing about this technology. Although it looks as if it must feel isolating to strap on a headset that shuts out the world around you, it could be great for socializing. When Facebook CEO Mark Zuckerberg spent $3 billion to buy the VR headset maker Oculus back in 2014, he pointed to the potential for social interactions as a key reason. And he's right. Virtual reality can give you a sense of being with others that a FaceTime call on an iPhone will never match. Virtual reality could make it more emotionally fulfilling to connect with far-flung friends and family, or lend a feeling of physical presence to online classes that is impossible to attain otherwise.

Virtual reality may look isolating, but the technology has untapped potential as a platform for socializing with people anywhere in the world.

But oddly enough, the social network hasn't made *social* applications a focal point for the Oculus Rift headset, which launched in 2016. Primarily it's still meant for playing games and watching short films. In April, Facebook released an app for Oculus Rift called Spaces to let you to get together with your Facebook friends in VR. But it's pretty humdrum. You can take a virtual selfie with your buddies, create a customized avatar based on your Facebook photos, and watch 360° videos or make 3-D doodles with a giant marker. You can interact only with the friends you already have on the social network, so unless your friends happen to have virtual-reality headsets too, you really can't do much in Spaces but hang out all by yourself. And even if you do meet up with a friend in Spaces, it gets old quickly: you're forced to stand around a virtual table the whole time. It seems that in order to avoid overwhelming early VR users, Facebook has gone too far in the direction of simplicity.

Rec Room

Altspace VR

High Fidelity

Facebook Spaces

Recently I found a virtual social place that's actually fun. It's a free app for Rift and another high-end headset, HTC's Vive, called Rec Room. Its virtual world is laid out like a cartoon version of a high school gym where you can play games like paintball and dodgeball, which you control with physical movements in real life. There's also a big communal lobby called the locker room (for hanging out, not for changing in and out of virtual clothes—you can do that in a private Rec Room dorm room). In the locker room you can meet up with friends or strangers, shoot hoops, or play Ping-Pong.

Rec Room has plenty of flaws, but it nonetheless shows the power of today's truly immersive virtual-reality technology to promote connections between people in ways that past attempts at virtual socializing—remember Second Life?—could never muster. The interactions with others are largely intuitive; to become friends with people in Rec Room, for instance, you shake their

hands, which produces buzzing feedback in the handheld controller. I've had a blast spending time in Rec Room with my one other friend who uses VR, who in real life lives across the country. And it's also the only virtual environment I've found that prompts you to connect with people you don't know in ways that aren't so awkward you want to rip off your headset.

Although market researcher IDC believes 10 million virtual-reality headsets shipped last year, that number is tiny compared with, say, the smartphone market, where 1.5 billion handsets were shipped in 2016. And I think the technology will struggle to snag more users—and thus to come down in price from at least $800 for a headset and VR-ready PC—unless it becomes more social. It's not that virtual reality isn't fun on your own. It's delightful to get transported somewhere by pulling a headset over your eyes and headphones over your ears. But no matter how entrancing virtual reality is, it's ultimately a lonely escape if nobody else is around to enjoy it with. If experiences like Rec Room catch on, VR may become the first true post-smartphone social platform.

Structured social

Rec Room, made by a company called Against Gravity, is not very advanced visually. When I pop in to the app, I embody what I think of as Geometric Rachel: an ovoid head, rectangular torso, and mitten-like hands, with glasses and a hairdo that looks like a bob-shaped helmet. But even within the confines of this simple avatar, I have a surprising amount of fun hiding in shipping containers and shooting at opponents (including my VR buddy) in a giant game of capture the flag.

One day I connected with Against Gravity cofounder and chief creative officer Cameron Brown for an interview in the app's locker room. His avatar was decked out in square purple glasses, a black baseball cap, and a white tank top with an orange Rec Room logo—a simple happy face—on the front. His hand suddenly fell to the ground; the app was having network issues, he said. He whisked me away from the clamor of the locker room to a big, empty lounge with a wood floor, a Ping-Pong table, and couches that our legless bodies couldn't actually sit on. It's the rare private space in Rec Room, which anyone can use and invite others into. He said some people have taken to using the space for meetings, and there are a conference table and whiteboards in one corner just for that.

"Even if you're just ... throwing darts into other people's faces, you've kind of got a reason to be there."

Since VR technology is still so new that most of us don't have lots of friends using it, and interacting with strangers in virtual reality is about as awkward as doing so in real life, Rec Room is using slightly unconventional methods to get you to socialize. It literally forces you into games. Brown calls it "structured social." If you and a few others enter the gym around the same time to play dodgeball, for instance, Rec Room's software will automatically sort everyone into teams and an overly chipper, British-accented woman's voice will suddenly start counting down the seconds until the match begins. Before you can think about it, you're thrust into a game. "Who enjoys going to a party where you don't know anybody, and just walking into a room and going, 'Hi'?" he said.

True enough. Still, Rec Room's ice-breaking methods can be jarring at first, like the time I suddenly found myself wielding a giant pen in a game of 3-D charades. Like most people, I do prefer to socialize with those whose interests I share, rather than complete strangers, and that's not necessarily Rec Room's aim (although you can set up private games with just your friends). Eventually, though, this sink-or-socialize tactic got me to loosen up and enjoy myself rather than stand around trying to work up the nerve to talk to others. "Even if you're just mucking around in the locker room throwing darts into other people's faces, you've kind of got a reason to be there, and before you know it, half an hour's gone by," Brown said.

Mindless fun or not, Rec Room was remarkable for how it gave me a sense of being in a space with other people, especially my friend who is thousands of miles away. We were giddy at being able to communicate and be together in a new way. It was much more enjoyable than Facebook Spaces, where there's not enough to do, and another social app, AltspaceVR, where there are too many options for socializing, none very entertaining.

Avatar issues

If you're hanging out in virtual reality, you're going to need a body, and what this body must look like, or whether it even has to be human, depends on the context. Often, it seems cartoonish human figures are best for staying clear of the uncanny valley, since it's still difficult to make avatars look just like us. (Probably the most realistic-looking avatar I've yet seen belongs to Second Life creator Philip Rosedale in his new social VR app, High Fidelity, and even that one doesn't look quite like him.)

Regardless of how well designed your avatar is in VR, one way these worlds resemble real life is that your perceived gender shapes the interactions you have. In Rec Room and other socially geared apps, like AltspaceVR and Facebook Spaces, I prefer to make my avatar



DAVID BRANDON GEE TING

Past attempts at creating social environments in VR weren't very compelling. But now that the technology is much richer and more immersive, it's more appealing to experience it with friends.

female—and preferably similar in appearance to me, with brown hair and, when it's an option, glasses. Being true to your actual identity can make you feel that your virtual self is authentic, but as a female character you're likely to face behavior that is obnoxious or worse.

I haven't experienced anything that bad. The worst was when some whiny, mouth-breathing adolescent boy complained that I sucked at a game: admittedly, it pissed me off and hurt my feelings a bit. (To be fair, I was terrible at several of the games.) But Brown acknowledges that harassment can be an issue. He told me of an incident in which the fiancée of a regular Rec Room user tried out the app and found a bunch of other players instantly in her face when she entered the locker room.

Simple tools to counter these behaviors are critical, but like VR itself, they're still very much in the experimental stage. Rec Room and AltspaceVR let you mute offensive users or build an invisible cushion around your avatar to make bullies disappear if they get too close. Rec Room also lets players vote on whether someone should be kicked out of a game. These tools will have to get more advanced and intuitive.

It could take several years for headset prices to drop and for a range of apps to figure out the interactions we want to have with each other in virtual spaces, making VR so compelling that, like the TV, PC, and smartphone, it becomes something the average person wants to use. Rec Room is not going to be that killer app for many people. But because it's focused on certain activities and thoughtful about how to promote social interactions, it's a good model for what's possible. I hope it can stick around until there are many more people to play with.

Rachel Metz is a senior editor at MIT Technology Review.



The Unaffordable Urban Paradise

Tech startups helped turn a handful of metro areas into megastars. Now they're tearing those cities apart.

By Richard Florida

LEO ESPINOSA



In the 1980s, I was part of a team doing research into the geography of the high-tech industry. We couldn't find a single significant high-tech company in an urban neighborhood. Instead, they were all out in the suburbs—not just Intel and Apple in Silicon Valley, or Microsoft in the Seattle suburbs, but the Route 128 beltway outside Boston, and the corporate campuses of North Carolina's Research Triangle.

Now everything's changed. In 2016, the San Francisco metro area was the top location for venture capital investment in the country, hauling in $23.4 billion—more than triple the VC investment in Silicon Valley proper. New York had virtually zero VC-backed startups in the 1980s, but

The tech industry's migration out of the suburbs and into cities like San Francisco, New York, and Boston.

last year it took in $7.6 billion, eclipsing Silicon Valley as well. Boston and Cambridge were close, with $6 billion. Los Angeles drew $5.5 billion. The likes of Google, Apple, Microsoft, and Facebook continue to maintain suburban campuses, but more than half of venture-capital-financed startups are now in dense urban neighborhoods. Amazon's headquarters are in downtown Seattle, and Google has now taken over the old Port Authority building in Manhattan.

For years, economists and mayors believed that high-tech development would "lift all boats." The reality is a new phase of winner-take-all urbanism.

The migration of high-tech startups to cities is less of a reversal and more of a historical correction. In 2006, the venture capital icon Paul Graham said that for all its power, Silicon Valley had a great weakness. The high-tech "paradise" created during the 1950s and 1960s "is now one giant parking lot," he said. "San Francisco and Berkeley are great, but they're 40 miles away. Silicon Valley proper is soul-crushing suburban sprawl. It has fabulous weather, which makes it significantly better than the soul-crushing sprawl of most other American cities. But a competitor that managed to avoid sprawl would have real leverage."

And that's what happened. Urban areas provide the diversity, creative energy, cultural richness, vibrant street life, and openness to new ideas that attract startup talent. Their industrial and warehouse buildings also provide employees with flexible and reconfigurable work spaces. Cities and startups are a natural match.

For years, economists, mayors, and urbanists believed that high-tech development was an unalloyed good thing, and that more high-tech startups and more venture capital investment would "lift all boats." But the reality is that high-tech development has ushered in a new phase of what I call winner-take-all urbanism, where a relatively small number of metro areas, and a small number of neighborhoods within them, capture most of the benefits.

Middle-class neighborhoods have been hollowed out in the process. In 1970, roughly two-thirds of Americans lived in middle-class neighborhoods; today less than 40 percent of us do. The middle-class share of the population shrank in a whopping 203 out of 229 U.S. metro areas between 2000 and 2014. And places where the middle class is smallest include such superstar cities and tech hubs as New York, San Francisco, Boston, Los Angeles, Houston, and Washington, D.C.

Despite all this, it wouldn't make any sense to put the brakes on high-tech development. Doing so would only cut off a huge source of innovation and economic development. High-tech industry remains a major driver of economic progress and jobs, and it provides much-needed tax revenues that cities can use to address and mitigate the problems that come with financial success.

Middle-class neighborhoods have been hollowed out. The middle-class share of the population shrank in 203 out of 229 U.S. metro areas between 2000 and 2014.

But if high-tech development causes problems, and stopping it doesn't solve those problems, what comes next?

High-tech companies should—out of self-interest, if for no other reason—embrace a shift to a kind of urbanism that allows many more people, especially blue-collar and service workers, to share in the gains of urban development. The superstar cities they've helped create cannot survive when nurses, EMTs, teachers, police officers, and other service providers can no longer afford to live in them.

Here's how they can do it. First, they can work with cities to help build more housing, which would reduce housing prices. They can support efforts to liberalize outdated zoning and building codes to enable more housing construction, and invest in the development of more affordable housing for service and blue-collar workers.

Second, they can work for, support, and invest in the development of more and better public transit to connect outlying areas to booming cores and tech clusters where employment is—and to spur and generate denser real estate and business development around those stops and stations.

Third, they can engage the wider business community and government to upgrade the jobs of low-wage service workers—who now make up more than 45 percent of the national workforce—into higher-paying, family-supporting work.

This last idea might seem outlandish, but it's analogous to how the U.S. turned low-paying manufacturing jobs of the early 20th century into middle-class jobs in the 1950s and 1960s. At age 13, my father left school to work in a factory. It took nine people—his father, his mother, six brothers and sisters, and him—to generate enough income to support the family. But when he came back after serving in World War II and took up work in the same factory, his old job paid enough to support a wife and kids, buy a house, and put my brother and me through college.



ROBERT NICKELSBERG/GETTY

The tech startup boom has brought billions of dollars of venture capital into urban areas like San Francisco. It's also driven out the middle class and caused a wave of resentment, exemplified by the Google Bus protests of 2011.

His job had been transformed by New Deal policies that boosted the pay of manufacturing workers.

We can do this for service jobs today, and we can start by lifting the wage floor. If we were to set the minimum wage to 50 percent of the prevailing local median wage, it would vary from city to city, ranging from a high of around $15 an hour in San Jose and Washington, D.C., to roughly $14 in San Francisco, about $13 in Boston, New York, and Seattle, and around $9.50 in less expensive places such as Las Vegas, Louisville, Memphis, Nashville, New Orleans, Orlando, and San Antonio.

In the 1980s and 1990s, leading manufacturing companies worked closely with their suppliers to upgrade manufacturing jobs, pay blue-collar workers more, and engage them in teamwork and lean production—which led to higher productivity and performance. Paying better wages to service workers would do for companies in the service sector what it once did for manufacturing firms.

Despite its creative energy and innovative prowess, America's tech industry has generated a host of challenges for cities. It's time to put its tremendous resources, talent, and skills to work helping cities solve the urban crisis it helped to create.

Richard Florida is a university professor at the University of Toronto's Rotman School of Management and the cofounder of the Atlantic*'s CityLab. His books include* The Rise of the Creative Class *and the just-released* The New Urban Crisis.



Why Bad Things Happen to Clean-Energy Startups

New technologies for storing power from wind and solar farms will be key to a clean-energy future. But Aquion Energy's recent bankruptcy shows the market challenge of making that happen.

By James Temple

By all measures that usually matter, Aquion Energy should have succeeded.

The startup, which sold big batteries for renewable projects and power grids, was founded by Jay Whitacre, a materials science professor at Carnegie Mellon who previously developed batteries for Mars rovers at NASA. It raised nearly $200 million from prominent investors, including Bill Gates, and venture capitalists at Kleiner Perkins and Shell. Perhaps most important of all, the company entered the market with a clear-eyed awareness of earlier missteps by battery startups. It took pains to avoid the use of rare materials. It relied on repurposed manufacturing equipment. And it identified niche markets where it might gain a foothold.

But on March 8, after failing to raise additional funding, Aquion filed for bankruptcy protection, cut 80 percent of its staff, and halted manufacturing. It was the latest of several stumbles for venture-backed storage startups. EnerVault, which was developing what are known as flow batteries, put itself up for sale after failing to find additional investors in 2015. Later that year, liquid-metal battery startup Ambri laid off a quarter of its staff. Around that same time, LightSail Energy, which was struggling to develop technology to store energy as compressed air in carbon-fiber tanks, pivoted to selling its containers to natural-gas suppliers. Taken together, these struggles have deflated hopes for the emergence of affordable and practical grid storage anytime soon.

And that's a problem. Without cheap ways of storing excess energy generated from intermittent sources like wind and sun, there are limitations on how much these renewable sources of power can contribute to the grid's overall electricity generation—and, by extension, how much we can cut the greenhouse-gas emissions driving climate change. There are already days when California solar farms have to shut down because they're generating more power than the grid can use at a particular time. And yet the system still needs to run enough fossil-fuel backup plants to meet total demand whenever the sun dips behind clouds.

A year ago, *MIT Technology Review* ranked Aquion number five on its list of 50 Smartest Companies. It's hard to decipher the particulars of what went wrong, as the bankruptcy filing offers few additional details. Whitacre declined requests to speak until an auction is complete, but he made clear that he hopes the company or technology can proceed in some form after that point.

What is clear is that despite the compelling need for better grid storage technology, any startup today faces several

Aquion Energy's filing for Chapter 11 bankruptcy March 8, 2017

DANIEL ZENDER

daunting realities. First, the slowly developing market for advanced grid storage still isn't large, in part because the technologies are immature and expensive. Second, and more important in the immediate term, the price of existing technology in the form of lithium-ion batteries has dropped far faster than expected, narrowing the promised benefits of new approaches like Aquion's.

"Don't hold your breath for the things that come after lithium-ion," says Ilan Gur, founding director of the Cyclotron Road program for energy entrepreneurs, who previously cofounded a battery company that was acquired by Bosch. "We're much more likely to ride the lithium-ion cost curve for another few decades."

By now it's a familiar phenomenon in clean energy. Companies must make a massive up-front investment to develop new hardware and scale up manufacturing, all while chasing moving price and performance targets as incumbent technologies improve. Faced with such challenges, vanishingly few succeed.

The danger, in this case, is that many observers believe lithium-ion isn't the right technology for full-scale baseload grid storage, because there seem to be limits on how cheap and long-lasting the technology can ever become. Yet its success frosts over an already chilly investment market for technologies that may be only minimally better today, but could have far greater potential to transform the energy system in the long run.

How can we support such technologies long enough to get to that point?

Toast

When *MIT Technology Review* first wrote about Aquion, in 2012, the company said it hoped to manufacture its batteries for less than $300 per kilowatt-hour of capacity. The technology paired a saltwater electrolyte with a manganese oxide cathode and a carbon-based anode. The price positioned the product between low-end lead-acid batteries and costlier lithium-ion technology. But a big part of the company's pitch was that the batteries lasted far longer and performed better in the heavy charging and discharging cycles inherent to the electric grid. That promised to lower the lifetime costs of the system.

It is still anyone's bet what advanced grid storage technology, or combination of different approaches, could ultimately displace lithium-ion batteries.

It's unclear where Aquion's pricing stood at the time of its bankruptcy filing, but lithium-ion batteries themselves fell below the $300-per-kilowatt-hour threshold last year, according to Bloomberg New Energy Finance. Specifically, the price dropped by nearly half in just two years, reaching $273 in 2016, as worldwide production ramped up to meet growing demand for phones, electric cars, and solar-power backup systems. And prices for lithium-ion batteries are likely to continue declining. Bloomberg projects they'll fall to $109 per kilowatt-hour by 2025, and $73 by 2030.

Donald Sadoway, a materials scientist at MIT and cofounder of Ambri, is skeptical that such estimates reflect total costs for lithium-ion storage systems at grid scale. He also argues that future projections from Bloomberg New Energy Finance are unrealistic, noting that the 2030 estimate seems to dip below the cost of raw materials. But either way, the sudden drop in price has already had a "huge ripple effect" on the storage market, he says. "When lithium-ion comes out with these audacious claims, it makes investors pause," Sadoway says. "They look and say, 'Wow, you guys are toast.'"

Meanwhile, customers considering a multimillion-dollar storage system have little incentive to bet on an emerging and riskier technology. Lithium-ion batteries already meet many of the specific needs of large-scale utility customers, offering a reliable product from stable providers, says John Zahurancik, president of the energy storage division of AES, a power project developer increasingly focused on battery systems. When California utilities were forced to quickly replace a natural-gas storage site after a major leak in 2015, they opted for three separate lithium-ion systems, assembled by AES and Tesla. "We see it as a technology that's here, now, and ready," Zahurancik says.

Charging ahead

But there are good reasons to hedge our battery bets, says Jessika Trancik, an associate professor in energy studies at MIT. As much as prices have dropped, lithium-ion batteries still aren't cheap enough to serve bigger roles in the grid. And it's not clear whether they'll get there, she says.

Pricing is also more complicated than what a single kilowatt-hour captures. The more telling metric is dollars per kilowatt-hour per cycle, a figure that would reflect lifetime costs of the system. And that gets us to a big problem for lithium-ion: these batteries can fade fast, as every smartphone owner knows. That makes them less than ideal for the strains of the baseload storage required to balance out the growing portion of intermittent sources like wind and solar in the grid.

It's still anyone's bet what grid storage technology, or combination of them, could ultimately displace lithium-ion batteries. In addition to advanced battery possibilities, there are flywheels, compressed air, hydrogen fuel cells, electric-vehicle-to-grid systems, and even energy-storing air conditioners. (Pumping water up a hill and letting it run down later remains by far the biggest storage technology.) But all

these technologies suffer from significant shortcomings, and all are likely to require deep sustained investments to further develop them, test whether they are truly viable, and make them cost competitive.

The question is where such investments will come from. The Trump administration has taken steps to dismantle the federal funding program for clean-energy startups, and venture capital investments for such technologies have dropped nearly 30 percent since 2011, from $7.5 billion to $5.2 billion, according to a recent report from the Brookings Institution.

The good news is that the market for large storage systems is widening as more wind and solar energy projects are built, and as factories look to shave costs during peak usage times. And the economics of grid storage are becoming more favorable as additional renewable generation comes online and more aging fossil-fuel plants shut down, say Severin Borenstein, an energy economist at the University of California, Berkeley. "As we roll out more intermittent renewables, storage becomes more valuable," he says.

The fear, however, is that market mechanisms alone aren't working fast enough given the pace at which the world needs to cut emissions to avoid the greatest threats of climate change. One tool that could accelerate the shift is smart public policy. Carbon taxes, or cap-and-trade programs like the ones in place in California, New York, and other states, raise the cost of using fossil-fuel plants and increase the incentives to add more renewables and storage to the grid. There are more direct ways the government can support the sector as well, such as California's requirement that state utilities add more than 1.3 gigawatts of storage to the grid by the end of this decade, or federal grant and loan programs for promising startups.

But ultimately, nascent storage technologies are fighting forces embedded deep in the core of capitalism. Markets consolidate around dominant technologies and companies. It often takes a radical advance to shake up the old order, and in the energy sector those don't come along often.

James Temple is senior editor for energy at MIT Technology Review.

Lead from **experience.**
Lead for **impact.**

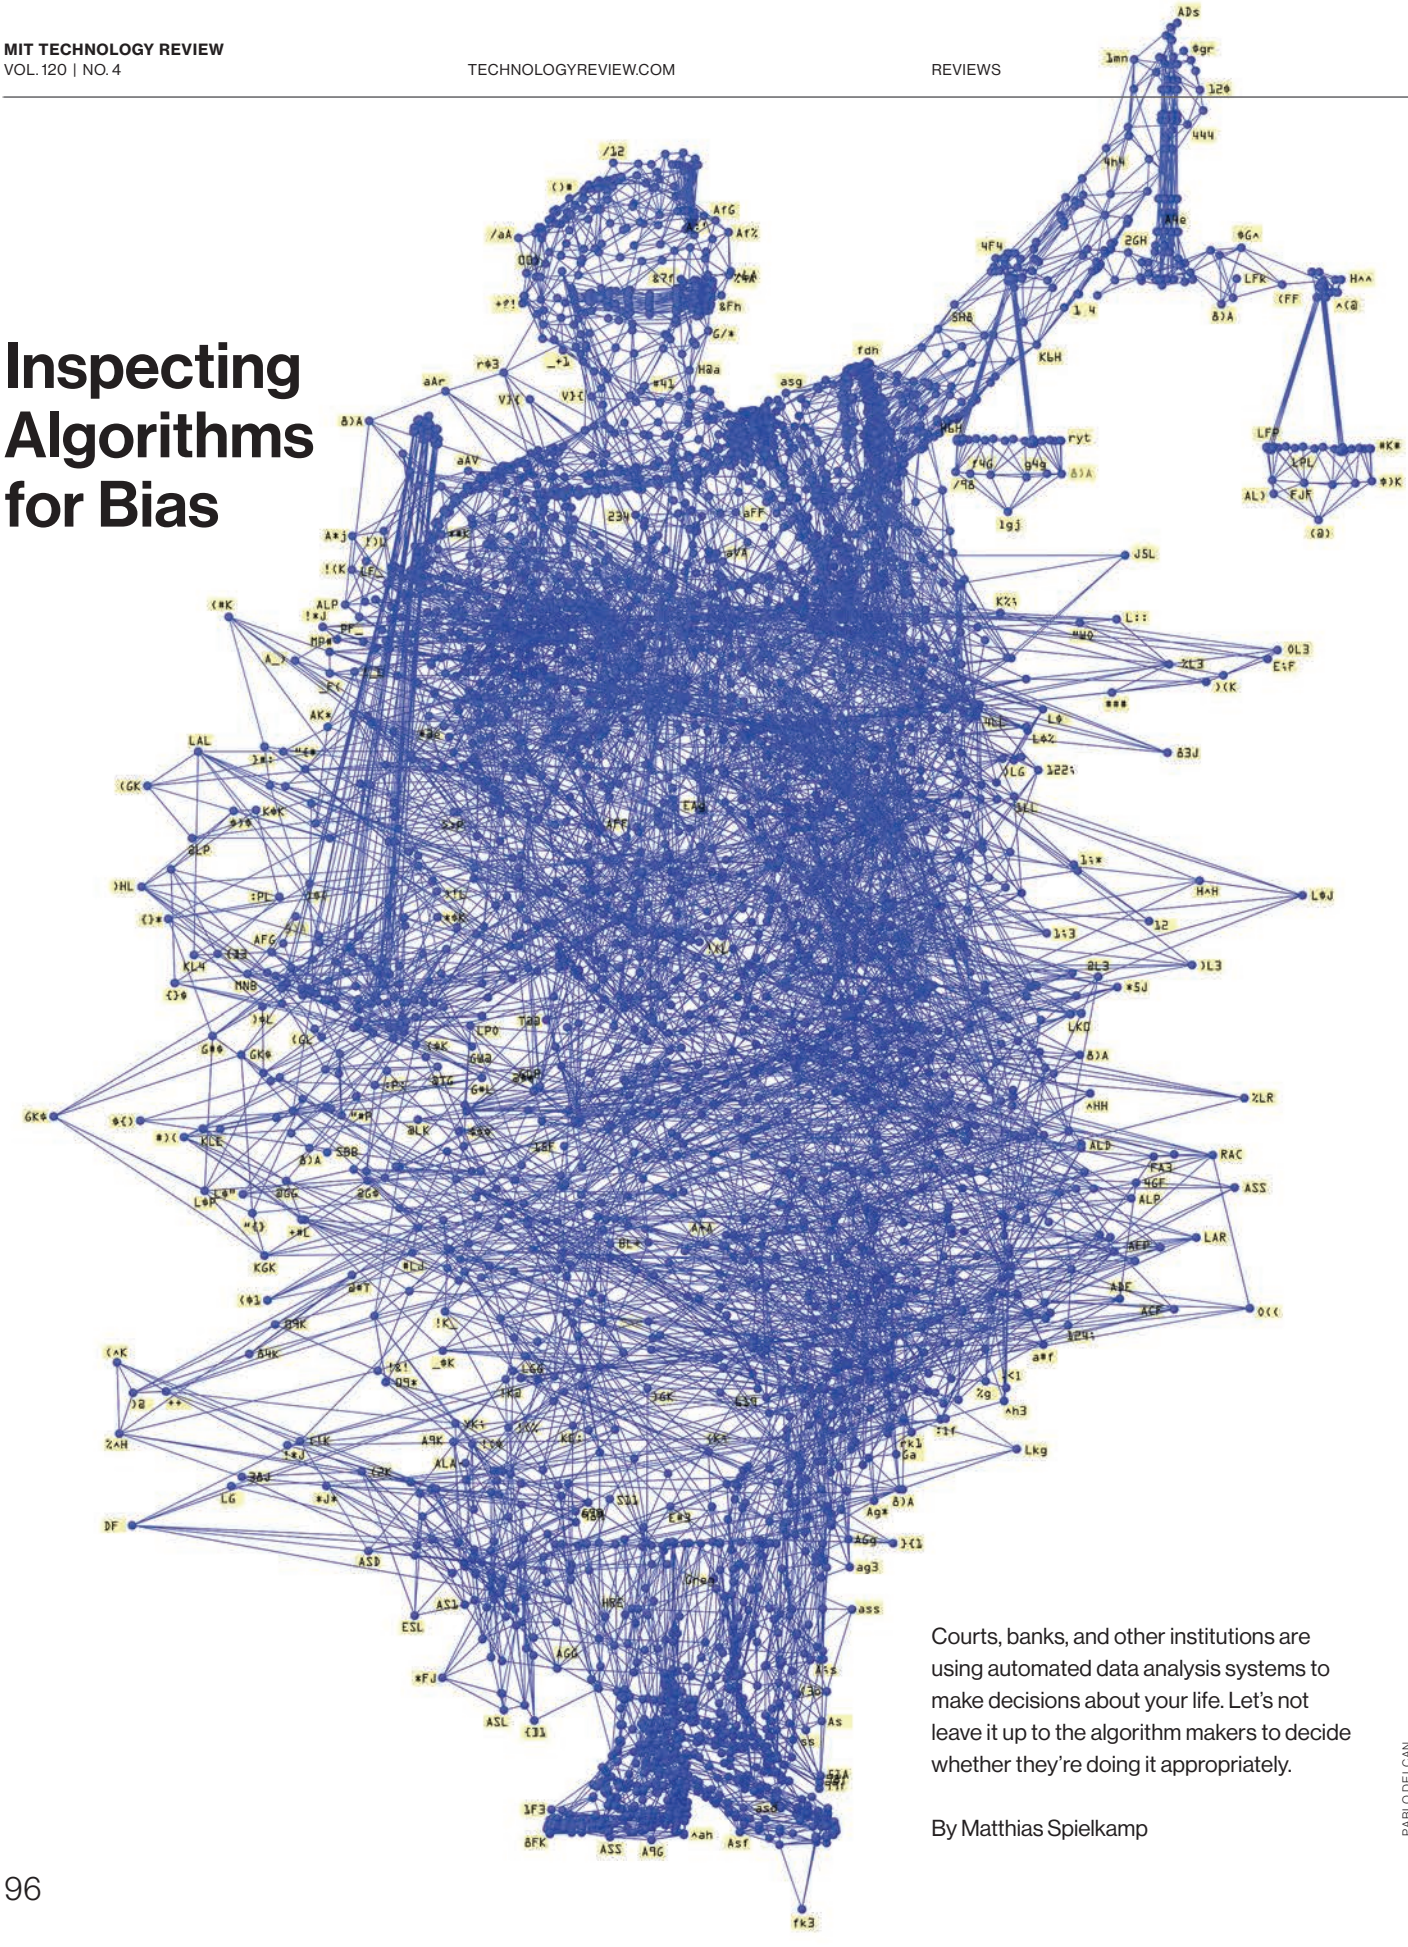
The **Harvard Advanced Leadership Initiative** offers a calendar year of rigorous education and reflection for highly accomplished leaders in business, government, law, medicine, technology and other sectors who are transitioning from their primary careers to their next years of service. Led by award-winning faculty members from across Harvard, the program aims to deploy a new leadership force tackling the world's most challenging social and environmental problems.





be inspired at **advancedleadership.harvard.edu**
+1-617-496-5479

Inspecting Algorithms for Bias



Courts, banks, and other institutions are using automated data analysis systems to make decisions about your life. Let's not leave it up to the algorithm makers to decide whether they're doing it appropriately.

By Matthias Spielkamp

PABLO DELCAN

It was a striking story. "Machine Bias," the headline read, and the teaser proclaimed: "There's software used across the country to predict future criminals. And it's biased against blacks."

ProPublica, a Pulitzer Prize–winning nonprofit news organization, had analyzed risk assessment software known as COMPAS. It is being used to forecast which criminals are most likely to reoffend. Guided by such forecasts, judges in courtrooms throughout the United States make decisions about the future of defendants and convicts, determining everything from bail amounts to sentences. When ProPublica compared COMPAS's risk assessments for more than 10,000 people arrested in one Florida county with how often those people actually went on to reoffend, it discovered that the algorithm "correctly predicted recidivism for black and white defendants at roughly the same rate." But when the algorithm was wrong, it was wrong in different ways for blacks and whites. Specifically, "blacks are almost twice as likely as whites to be labeled a higher risk but not actually re-offend." And COMPAS tended to make the opposite mistake with whites: "They are much more likely than blacks to be labeled lower risk but go on to commit other crimes."

Whether it's appropriate to use systems like COMPAS is a question that goes beyond racial bias. The U.S. Supreme Court might soon take up the case of a Wisconsin convict who says his right to due process was violated when the judge who sentenced him consulted COMPAS, because the workings of the system were opaque to the defendant. Potential problems with other automated decision-making (ADM) systems exist outside the justice system, too. On the basis of online personality tests, ADMs are helping to determine whether someone is the right person for a job. Credit-scoring algorithms play an enormous role in whether you get a mortgage, a credit card, or even the most cost-effective cell-phone deals.

"Machine Bias"
ProPublica, May 23, 2016

"COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity"
Northpointe, July 8, 2016

"Technical Response to Northpointe"
ProPublica, July 29, 2016

"False Positives, False Negatives, and False Analyses: A Rejoinder to 'Machine Bias'"
Anthony Flores, Christopher Lowenkamp, and Kristin Bechtel
August 10, 2016

Lawmakers, the courts, and an informed public should decide what we want algorithms to prioritize.

It's not necessarily a bad idea to use risk assessment systems like COMPAS. In many cases, ADM systems can increase fairness. Human decision making is at times so incoherent that it needs oversight to bring it in line with our standards of justice. As one specifically unsettling study showed, parole boards were more likely to free convicts if the judges had just had a meal break. This probably had never occurred to the judges. An ADM system could discover such inconsistencies and improve the process.

But often we don't know enough about how ADM systems work to know whether they are fairer than humans would be on their own. In part because the systems make choices on the basis of underlying assumptions that are not clear even to the systems' designers, it's not necessarily possible to determine which algorithms are biased and which ones are not. And even when the answer seems clear, as in ProPublica's findings on COMPAS, the truth is sometimes more complicated.

What should we do to get a better handle on ADMs? Democratic societies need more oversight over such systems than they have now. AlgorithmWatch, a Berlin-based nonprofit advocacy organization that I cofounded with a computer scientist, a legal philosopher, and a fellow journalist, aims to help people understand the effects of such systems. "The fact that most ADM procedures are black boxes to the people affected by them is not a law of nature. It must end," we assert in our manifesto. Still, our take on the issue is different from many critics'—because our fear is that the technology could be demonized undeservedly. What's important is that societies, and not only algorithm makers, make the value judgments that go into ADMs.

Measures of fairness

COMPAS determines its risk scores from answers to a questionnaire that explores a defendant's criminal history and attitudes about crime. Does this produce biased results?

After ProPublica's investigation, Northpointe, the company that developed COMPAS, disputed the story, arguing that the journalists misinterpreted the data. So did three criminal-justice researchers, including one from a justice-reform organization. Who's right—the reporters or the researchers? Krishna Gummadi, head of the Networked Systems Research Group at the Max Planck Institute for Software Systems in Saarbrücken, Germany, offers a surprising answer: they all are.

Gummadi, who has extensively researched fairness in algorithms, says ProPublica's and Northpointe's results don't contradict each other. They differ because they use different measures of fairness.

Imagine you are designing a system to predict which criminals will reoffend. One option is to optimize for "true positives," meaning that you will identify as many people as possible who are at high risk of committing another crime. One problem with this approach is that it tends to increase the number of false positives: people who will be unjustly classified as likely reoffenders. The dial can be adjusted to deliver as few false positives as possible, but that tends to create more false negatives: likely reoffenders who slip through and get a more lenient treatment than warranted.

Raising the incidence of true positives or lowering the false positives are both ways to improve a statistical measure known as positive predictive value, or PPV. That is the percentage of all positives that are true.

As Gummadi points out, ProPublica compared false positive rates and false negative rates for blacks and whites and found them to be skewed in favor of whites. Northpointe, in contrast, compared the PPVs for different races and found them to be similar. In part because the recidivism rates for blacks and whites do in fact differ, it is mathematically likely that the positive predictive values for people in each group will be similar while the rates of false negatives are not.

One thing this tells us is that the broader society—lawmakers, the courts, an informed public—should decide what we want such algorithms to prioritize. Are we primarily interested in taking as few chances as possible that someone will skip bail or reoffend? What trade-offs should we make to ensure justice and lower the massive social costs of imprisonment?

If used properly, criminal-justice algorithms offer "the chance of a generation, and perhaps a lifetime, to reform sentencing and unwind mass incarceration in a scientific way."

No matter which way the dials are set, any algorithm will have biases—it is, after all, making a prediction based on generalized statistics, not on someone's individual situation. But we can still use such systems to guide decisions that are wiser and fairer than the ones humans tend to make on their own.

The controversy surrounding the New York Police Department's stop-and-frisk practices helps to show why. Between January 2004 and June 2012, New York City police conducted 4.4 million stops under a program that allowed officers to temporarily detain, question, and search people on the street for weapons and other contraband. But in fact, "88 percent of the 4.4 million stops resulted in no further action—meaning a vast majority of those stopped were doing nothing wrong," the *New York Times* said in an editorial decrying the practice. What's more: "In about 83 percent of cases, the person stopped was black or Hispanic, even though the two groups accounted for just over half the population." This example of human bias, illuminated through data analysis, is a reminder that ADM systems could play a positive role in criminal justice. Used properly, they offer "the chance of a generation, and perhaps a lifetime, to reform sentencing and unwind mass incarceration in a scientific way," according to Anthony Flores, Christopher Lowenkamp, and Kristin Bechtel, three researchers who found flaws in the methodology that ProPublica used to analyze COMPAS. The authors worry that this opportunity "is slipping away because of misinformation and misunderstanding" about the technology.

But if we accept that algorithms might make life fairer if they are well designed, how can we know whether they are so designed?

Democratic societies should be working now to determine how much transparency they expect from ADM systems. Do we need new regulations of the software to ensure it can be properly inspected? Lawmakers, judges, and the public should have a say in which measures of fairness get prioritized by algorithms. But if the algorithms don't actually reflect these value judgments, who will be held accountable?

These are the hard questions we need to answer if we expect to benefit from advances in algorithmic technology.

Matthias Spielkamp is executive director of AlgorithmWatch, an advocacy group that analyzes the risks and opportunities of automated decision making.





Access the future
with the authority on
cutting-edge technology

MIT
Technology
Review

SUBSCRIBE TODAY
technologyreview.com/julysubscribe

40 Years Ago



The Inevitable EV

Four decades ago, the electric car looked set to be an acceptable commuting vehicle within a few years.

"The electric car seems inevitable. As oil resources dwindle, it will appear far more prudent to develop an automobile powered by electricity, charged by an already-in-place power generation system, than to invest in huge plants to produce synthetic gasoline. Electric automobiles would snuggle nicely into a large, unexploited niche in the electric economy by recharging at night when electric demand is low and power plant capacity idle.

But as yet electric cars have proven abysmal technical failures. For example, in its 1975 tests of two electric cars on the market, *Consumer Reports* magazine found 'major safety and operating problems.' The cars performed poorly, not even living up to their modestly advertised abilities. They accelerated sluggishly, balked at hills, and petered out at inopportune times. The magazine declared that 'it would be foolhardy to drive either car on any public road.' The electric cars now on the market have a range of only about 30 miles at 50 miles per hour, with 0-to-30 acceleration times of about 12 seconds. What's more, the cars were not significantly cheaper to run than comparable gasoline-powered subcompacts.

Two test electric vehicles recently commissioned by the Energy Research and Development Administration could mark some resurgence for the electric car. The two electric car versions are being built for E.R.D.A. by General Electric Co. and AiResearch Manufacturing Co.

Both vehicles will possess about the same range and performance characteristics—70 miles between charges, 60-mile-an-hour speeds, and 0-to-30 m.p.h. in about eight seconds. They will also both be about the size and weight of a conventional subcompact, and will operate for about the same costs as a conventional car.

Officials at E.R.D.A. expect the two cars' performance to help the electric car become an acceptable commuting vehicle in a few years, even using today's relatively low-efficiency lead-acid batteries. The high-energy batteries expected in about a decade should improve performance even further.

Even the E.R.D.A. electric cars must overcome many safety and operating obstacles, however. For example, *Consumer Reports* cited the hazards of sulfuric acid spilling from ruptured batteries in the event of an accident. And the magazine's engineers discovered that during cold weather battery performance plummeted, rendering the electric cars they tested practically useless."

Excerpted from "Zapmobile," in the July/August 1977 issue of Technology Review.

MIT Technology Review (ISSN 1099-274X), July/August 2017 Issue, Reg. U.S. Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2017. The editors seek diverse views, and authors' opinions do not represent the official policies of their institutions or those of MIT. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the Internet at www.technologyreview.com/customerservice. Basic subscription rates: $41.94 per year within the United States; in all other countries, US$54. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051 Fort Erie, ON L2A 6C7. Printed in U.S.A. Audited by the Alliance for Audited Media

EVERY 9.8 SECONDS YOU REACH FOR US.

BUT YOU STILL DON'T KNOW US.

Qualcomm created technology that enables your smartphone to do all those things you love. From download speeds, to video, to voice recognition, to faster battery charge, to longer battery life, to security, to GPS, none of it would work the way you count on without Qualcomm getting there first.

qualcomm.com/weinvent

QUALCOMM®

Why you love your smartphone.

© 2017 Qualcomm Technologies, Inc. All Rights Reserved. Qualcomm is a trademark of Qualcomm Incorporated, registered in the United States and other countries.


Accelerating next
Hewlett Packard
Enterprise
> INTRODUCING THE WORLD'S MOST SECURE INDUSTRY STANDARD SERVERS
Protect your data like it's nobody's business. Because it isn't. Experience unique protection from server silicon to software – so you'll never have to compromise. It's all part of the new generation compute experience from HPE.
Learn more at hpe.com/security
©Copyright 2017 Hewlett Packard Enterprise Development LP.